Vertex Cover Reconfiguration and Beyond

Amer E. Mouawad[1⋆], Naomi Nishimura[1⋆], and Venkatesh Raman[2]

[1] David R. Cheriton School of Computer Science
University of Waterloo, Ontario, Canada.
{aabdomou, nishi}@uwaterloo.ca
[2] The Institute of Mathematical Sciences
Chennai, India.
vraman@imsc.res.in

arXiv:1402.4926v1 [cs.CC] 20 Feb 2014

Abstract. In the Vertex Cover Reconfiguration (VCR) problem, given graph $G = (V, E)$, positive integers k and ℓ, and two vertex covers S and T of G of size at most k, we determine whether S can be transformed into T by a sequence of at most ℓ vertex additions or removals such that every operation results in a vertex cover of size at most k. Motivated by recent results establishing the $\mathbf{W}[\mathbf{1}]$-hardness of VCR when parameterized by ℓ, we delineate the complexity of the problem restricted to various graph classes. In particular, we show that VCR remains $\mathbf{W}[\mathbf{1}]$-hard on bipartite graphs, is **NP**-hard but fixed-parameter tractable on graphs of bounded degree, and is solvable in time polynomial in $|V(G)|$ on cactus graphs. We prove $\mathbf{W}[\mathbf{1}]$-hardness and fixed-parameter tractability via two new problems of independent interest.

1 Introduction

Under the reconfiguration framework, we consider structural and algorithmic questions related to the solution space of a search problem $\mathcal{Q}$. Given an instance $\mathcal{I}$, an optional range $[r_l, r_u]$ bounding a numerically quantifiable property Ψ of feasible solutions for $\mathcal{Q}$, and a symmetric adjacency relation (usually polynomially-testable) $\mathcal{A}$ on the set of feasible solutions, we can construct a *reconfiguration graph* $R_{\mathcal{Q}}(\mathcal{I}, r_l, r_u)$ for each instance $\mathcal{I}$ of $\mathcal{Q}$. The nodes of $R_{\mathcal{Q}}(\mathcal{I}, r_l, r_u)$ correspond to the feasible solutions of $\mathcal{Q}$ having $r_l \leq \Psi \leq r_u$, and there is an edge between two nodes whenever the corresponding solutions are adjacent under $\mathcal{A}$. An edge can be seen as a *reconfiguration step* transforming one solution into the other. Given two feasible solutions for $\mathcal{I}$, S and T, one can ask if there exists a walk (*reconfiguration sequence*) in $R_{\mathcal{Q}}(\mathcal{I}, r_l, r_u)$ from S to T, or for the shortest such walk. On the structural side, one could ask about the diameter of reconfiguration graph $R_{\mathcal{Q}}(\mathcal{I}, r_l, r_u)$ or whether it is connected with respect to some or any $\mathcal{I}$, fixed $\mathcal{A}$, and fixed Ψ.

These types of reconfiguration questions have received considerable attention in recent literature [10, 12, 15, 17, 18] and are interesting for a variety of reasons. From an algorithmic standpoint, reconfiguration problems model dynamic situations in which we seek to transform a solution into a more desirable one, maintaining feasibility during the process. Reconfiguration also models questions of evolution; it can represent the evolution of a genotype where only individual mutations are allowed and all genotypes must satisfy a certain fitness threshold, i.e. be feasible. Moreover, the study of reconfiguration yields insights into the structure of the solution space of the underlying problem, crucial for the design of efficient algorithms. In fact, one of the initial motivations behind such questions was to study the performance of heuristics [12] and random sampling methods [3], where connectivity and other properties of the solution space play a crucial role.

⋆ Research supported by the Natural Science and Engineering Research Council of Canada.

Reconfiguration problems have so far been studied mainly under classical complexity assumptions, with most work devoted to determining the existence of a reconfiguration sequence between two given solutions. For most **NP**-complete problems, this question has been shown to be **PSPACE**-complete [15, 16, 19], while for some problems in **P**, the reconfiguration question could be either in **P** [15] or **PSPACE**-complete [2]. As **PSPACE**-completeness implies that the number of vertices in reconfiguration graphs, and therefore the length of reconfiguration sequences, can be exponential in the number of vertices in the input graph, it is natural to ask whether we can achieve tractability if we allow the running time to depend on the length of the sequence or on other properties of the problem. These results motivated Mouawad et al. [21] to study reconfiguration under the *parameterized complexity* framework [6].

Very recently [21], the Vertex Cover Reconfiguration (VCR) problem was shown to be fixed-parameter tractable when parameterized by k and $\mathbf{W}[\mathbf{1}]$-hard when parameterized by ℓ; in $R_{\mathrm{VC}}(G, 0, k)$ each feasible solution for instance G is a vertex cover of size at most k (a subset $S \subseteq V(G)$ such that each edge in E has at least one endpoint in S) and two solutions are adjacent if one can be obtained from the other by the addition or removal of a single vertex of G. Motivated by these results, we embark on a systematic investigation of the parameterized complexity of the problem restricted to various graph classes.

In Section 3, we start by showing that the VCR problem parameterized by ℓ remains $\mathbf{W}[\mathbf{1}]$-hard when restricted to bipartite graphs. In doing so, we introduce the (k, d)-Bipartite Constrained Crown problem and show that it plays a central role for determining the complexity of the reconfiguration problem. As Vertex Cover is solvable in polynomial time on bipartite graphs, this result provides the first example of a search problem in **P** whose reconfiguration version is $\mathbf{W}[\mathbf{1}]$-hard parameterized by ℓ, answering a question left open by Mouawad et al. [21]. In Section 4, we characterize instances of the VCR problem solvable in time polynomial in $|V(G)|$, and apply this characterization to trees and cactus graphs. Finally, in Section 5, we present the first fixed-parameter algorithm for the VCR problem parameterized by ℓ on graphs of bounded degree after establishing the **NP**-hardness of the problem on 4-regular graphs and introducing the Vertex Cover Walk problem.

2 Preliminaries

For general graph theoretic definitions, we refer the reader to the book of Diestel [5]. Unless otherwise stated, we assume that each graph G is a simple, undirected graph with vertex set $V(G)$ and edge set $E(G)$, where $|V(G)| = n$ and $|E(G)| = m$. The *open neighborhood* of a vertex v is denoted by $N_G(v) = \{u \mid \{u, v\} \in E(G)\}$ and the *closed neighborhood* by $N_G[v] = N_G(v) \cup \{v\}$. For a set of vertices $S \subseteq V(G)$, we define $N_G(S) = \{v \notin S \mid \{u, v\} \in E(G), u \in S\}$ and $N_G[S] = N_G(S) \cup S$. The subgraph of G induced by S is denoted by $G[S]$, where $G[S]$ has vertex set S and edge set $\{(u, v) \in E(G) \mid u, v \in S\}$.

A *walk* of length k from v_0 to v_k in G is a vertex sequence $v_0, \ldots, v_k$, such that for all $i \in \{0, \ldots, k-1\}$, $\{v_i, v_{i+1}\} \in E(G)$. It is a *path* if all vertices are distinct, and a *cycle* if $k \geq 3$, $v_0 = v_k$, and $v_0, \ldots, v_{k-1}$ is a path. We say vertices s and t (vertex sets A and B) are *separated* if there is no edge $\{s, t\}$ (no edge $\{a, b\}$ for $a \in A$, $b \in B$). Two subgraphs of G are *separated* whenever their corresponding vertex sets are separated. The *diameter* of a connected graph G, $diam(G)$, is the maximum distance from s to t over all vertex pairs s and t. The diameter of a disconnected graph is the maximum of the diameters of its connected components. For $r \geq 0$, the

r-neighborhood of a vertex $v \in V(G)$ is defined as $N^r_G[v] = \{u \mid dist_G(u,v) \leq r\}$. We also use the notation $B(v,r) = N^r_G[v]$ and call it a *ball of radius r around v*. For $A \subseteq V(G)$, we define $B(A,r) = \bigcup_{v \in A} N^r_G[v]$. A *matching* $\mathcal{M}(G)$ on a graph G is a set of edges of G such that no two edges share a vertex; we use $V(\mathcal{M}(G))$ to denote the set of vertices incident to edges in $\mathcal{M}(G)$. A set of vertices $A \subseteq V(G)$ is said to be *saturated* by $\mathcal{M}(G)$ if $A \subseteq V(\mathcal{M}(G))$.

To avoid confusion, we refer to *nodes* in reconfiguration graphs, as distinguished from *vertices* in the input graph. We denote an instance of the Vertex Cover Reconfiguration problem by (G, S, T, k, ℓ), where G is the input graph, S and T are the *source* and *target* vertex covers respectively, k is the *maximum allowed capacity*, and ℓ is an upper bound on the length of the reconfiguration sequence we seek in $R_{\text{VC}}(G, 0, k)$. By a slight abuse of notation, we use upper case letters to refer to both a node in the reconfiguration graph as well as the corresponding vertex cover. For any node $S \in V(R_{\text{VC}}(G, 0, k))$, the quantity $k - |S|$ corresponds to the *available capacity* at S. We partition $V(G)$ into the sets $C_{ST} = S \cap T$ (vertices common to S and T), $S_R = S \setminus C_{ST}$ (vertices to be removed from S in the course of reconfiguration), $T_A = T \setminus C_{ST}$ (vertices to be added to form T), and $O_{ST} = V(G) \setminus (S \cup T) = V(G) \setminus (C_{ST} \cup S_R \cup T_A)$ (all other vertices). To simplify notation, we sometimes use G_Δ to denote the graph induced by the vertices in the symmetric difference of S and T, i.e. $G_\Delta = G[S \Delta T] = G[S_R \cup T_A]$. We say a vertex is *touched* in the course of a reconfiguration sequence from S to T if v is either added or removed at least once. We say a vertex v, in a vertex cover S, is *removable* if and only if $v \in S$ and $N_G(v) \subseteq S$.

Observation 1 *For any graph G and any two vertex covers S and T of G, $G_\Delta = G[S_R \cup T_A]$ is bipartite. Moreover, there are no edges between vertices in $S_R \cup T_A$ and vertices in O_{ST}.*

Proof. None of the vertices in S_R are included in T. Since T is a vertex cover of G, each edge of G must have an endpoint in T, and hence $G[S_R]$ must be an independent set. Similar arguments apply to $G[T_A]$ and to show that there are no edges between vertices in $S_R \cup T_A$ and vertices in O_{ST}. □

Observation 2 *For a graph G and any two vertex covers S and T of G, any vertex in $S_R \cup T_A$ must be touched an odd number of times and any vertex not in $S_R \cup T_A$ must be touched an even number of times in any reconfiguration sequence of length at most ℓ from S to T. Moreover, any vertex can be touched at most $\ell - |S_R \cup T_A| + 1$ times.*

Throughout this work, we implicity consider the Vertex Cover Reconfiguration problem as a parameterized problem with ℓ as the parameter. The reader is referred to the excellent books of Downey, Fellows, Flum, Grohe, and Niedermeier for more on parameterized complexity [6, 9, 22].

2.1 Representing Reconfiguration Sequences

There are multiple ways of representing a reconfiguration sequence between two vertex covers of a graph G. In Sections 3 and 4, we focus on a representation which consists of an ordered sequence of vertex covers or nodes in the reconfiguration graph. Given a graph G and two vertex covers of G, A_0 and A_j, we denote a reconfiguration sequence from A_0 to A_j by $\alpha = (A_0, A_1, \ldots, A_j)$, where A_i is a vertex cover of G and A_i is obtained from A_{i-1} by either the removal or the addition of a single vertex from A_{i-1} for all $0 < i \leq j$. For any pair of consecutive vertex covers (A_{i-1}, A_i) in α, we say A_i (A_{i-1}) is the *successor* (*predecessor*) of A_{i-1} (A_i). A reconfiguration sequence $\beta = (A_0, A_1, \ldots, A_i)$ is a *prefix* of $\alpha = (A_0, A_1, \ldots, A_j)$ if $i < j$.

In Section 5, we use the notion of edit sequences, which are a generalization of reconfiguration sequences. Given a graph $G = (V, E)$, we assume all vertices of G are labeled from 1 to n, i.e. $V = \{v_1, v_2, \ldots, v_n\}$. We let $\mathcal{E}_a = \{a_1, \ldots, a_n\}$ and $\mathcal{E}_r = \{r_1, \ldots, r_n\}$ denote the sets of *labeled addition markers* and *labeled removal markers*, respectively. We let $\mathcal{E}_a \cup \{a\}$ and $\mathcal{E}_r \cup \{r\}$ denote the sets of *addition markers* and *removal markers*, respectively. An *edit sequence* α is an ordered sequence of elements obtained from the full set of *markers* $\mathcal{E} = \{\varnothing, a, r\} \cup \mathcal{E}_a \cup \mathcal{E}_r$, where $\varnothing$ stands for blank, a stands for add, r stands for remove, a_i stands for add vertex v_i, r_j stands for remove vertex v_j, and $1 \leq i, j \leq n$.

For any edit sequence α, we say α is *unlabeled* if it only contains elements from the set $\{\varnothing, a, r\}$, α is *partly labeled* if it contains at least one element from each of the sets $\{a, r\}$ and $\mathcal{E}_a \cup \mathcal{E}_r$, and α is *labeled* if it contains no elements from the set $\{a, r\}$. We say α is *partial* whenever it contains one or more occurrences of the element $\varnothing$ and is *full* otherwise.

Observation 3 *The total number of possible full (partial) unlabeled edit sequences of size at most ℓ is $\sum_{i=1}^{\ell} 2^i < 2^{\ell+1}$ ($\sum_{i=1}^{\ell} 3^i < 3^{\ell+1}$).*

The *size* or the *length* of α, $|\alpha|$, is equal to the total number of markers in α. We let $\alpha[p] \in \mathcal{E}$, $1 \leq p \leq |\alpha|$, denote the marker at position p in α. We say $\alpha[p]$ is a *blank marker* whenever $\alpha[p] = \varnothing$. We say β is a *segment* of α whenever β consists of a subsequence of α with no gaps. The *size* or the *length* of a segment is defined as the total number of markers it contains. We use the notation $\alpha[p_1, p_2]$, $1 \leq p_1, p_2 \leq |\alpha|$, to denote the segment starting at position p_1 and ending at position p_2. Two segments β and β' are *consecutive* whenever β' occurs later than β in α and there are no gaps between β and β'. A *piece* is a group of zero or more consecutive segments. For any pair of consecutive segments β and β' in α, we say β' (β) is the *successor* (*predecessor*) of β (β'). Given an edit sequence α, a segment β of α is an *add-remove segment* if β contains addition markers followed by removal markers. We say β is a *d-add-remove segment*, $d > 0$, if β is an add-remove segment containing 1 to d addition markers followed by 1 to d removal markers.

Definition 1. *Given a positive integer $d > 0$, an edit sequence α is* d-well-formed *if it is subdivided into three consecutive pieces such that:*

(1) The starting piece *consists of zero or more removal markers,*
(2) the central piece *consists of zero or more d-add-remove segments, and*
(3) the ending piece *consists of zero or more addition markers.*

We introduce some useful operations on edit sequences. The *concatenation* of two edit sequences β and γ, denoted by $concat(\beta, \gamma)$, results in an edit sequence α. The size of α is $|\beta| + |\gamma|$, $\alpha[p] = \beta[p]$ for all $1 \leq p \leq |\beta|$, and $\alpha[p] = \gamma[p - |\beta|]$ for all $|\beta| + 1 \leq p \leq |\beta| + |\gamma|$. Given an edit sequence β and two positive integers p_1 and p_2, $1 \leq p_1, p_2 \leq |\beta|$, the *cut* operation, denoted by $cut(\beta, p_1, p_2)$, results in a new edit sequence $\alpha = concat(\beta[1, p_1 - 1], \beta[p_2 + 1, |\alpha|])$. When $p_1 = 1$, we assume $\beta[1, p_1 - 1]$ corresponds to the empty sequence. We write $cut(\beta, p_1)$ whenever $p_1 = p_2$. The *clean* operation, denoted by $clean(\beta)$, results in a new edit sequence α which is obtained by repeatedly applying $cut(\beta, p)$ as long as $\beta[p] = \varnothing$. Given a partial edit sequence β and a full edit sequence γ, the *merging* operation consists of replacing the pth blank marker in β with the pth marker in γ. We say γ is a *filling edit sequence* of β if $merge(\beta, \gamma)$ produces α, where β is partial and α and γ are not. Hence, $|clean(\beta)| + |\gamma|$ must be equal to $|\alpha|$. Finally, given $t \geq 2$ full labeled edit sequences $\alpha_1, \ldots, \alpha_t$, the *mixing* of those sequences, $mix(\alpha_1, \ldots, \alpha_t)$, produces the set of all full labeled edit

sequences of size $|\alpha_1| + \ldots + |\alpha_t|$. Each $\alpha \in mix(\alpha_1, \ldots, \alpha_t)$ consists of all markers in each α_i, $1 \leq i \leq t$. However, the respective orderings of markers from each α_i is maintained, i.e. if we cut from α the markers of all sequences except α_1, we get α_1.

We now discuss how edit sequences relate to reconfiguration sequences. Given a graph G and an edit sequence α, we use $V(\alpha)$ to denote the set of vertices touched in α, i.e. $V(\alpha) = \{v_i \mid a_i \in \alpha \vee r_i \in \alpha\}$. When α is full and labeled, we let $V(S, \alpha)$ denote the set of vertices obtained after executing all reconfiguration steps in α on G starting from some vertex cover S of G. We say α is *valid* whenever every set $V(S, \alpha[1, p])$, $1 \leq p \leq |\alpha|$, is a vertex cover of G and we say α is *invalid* otherwise. Note that even if $|S| \leq k$, α is not necessarily a walk in $R_{\text{VC}}(G, 0, k)$, as α might violate the maximum allowed capacity constraint k. Hence, we let $cap(\alpha) = max_{1 \leq p \leq |\alpha|}(|V(S, \alpha[1, p])|)$ and we say α is *tight* whenever it is valid and $cap(\alpha) \leq k$. A partial labeled edit sequence α is valid or tight if $clean(\alpha)$ is valid or tight.

Observation 4 *Given a graph G and two vertex covers S and T of G, an edit sequence α is a reconfiguration sequence from S to T if and only if α is a tight edit sequence from S to T.*

Given a graph G, a vertex cover S of G, a full unlabeled edit sequence α, and an ordered sequence $L = \{l_1, \ldots, l_{|\alpha|}\}$ of (not necessarily distinct) labels between 1 and n, the *label* operation, denoted by $label(\alpha, L)$, returns a full labeled edit sequence α'; every marker in α' is copied from α and assigned the corresponding label from L. We say a full unlabeled edit sequence α can be *applied* to G and S if there exists an L such that $label(\alpha, L)$ is valid starting from S.

Definition 2. *Given a graph G, a vertex cover S of G, and $t \geq 2$ valid labeled edit sequences α_1, ..., α_t, we say α_1, ..., α_t are* compatible *if every $\alpha \in mix(\alpha_1, \ldots, \alpha_t)$ is a valid edit sequence starting from S. They are* incompatible *otherwise.*

Observation 5 *Given $t \geq 2$ valid labeled edit sequences α_1, ..., α_t, if $V(\alpha_i)$ and $V(\alpha_j)$ are separated, for all $1 \leq i, j \leq t$ and $i \neq j$, then α_1, ..., α_t must be compatible.*

On graphs of bounded degree, edit sequences have certain special properties, which we capture using the notion of nice edit sequences. Moreover, we present the ConvertToNice algorithm (Algorithm 1) which transforms any valid edit sequence α from S to T into a nice edit sequence α' from S to T, such that $|V(S, \alpha'[1, p])| \leq |V(S, \alpha[1, p])|$ for all $1 \leq p \leq |\alpha|$. Hence, if α is tight then so is α'. Moreover, this transformation can be accomplished in $\mathcal{O}(n^4|\alpha|^4 2^d)$ time.

Definition 3. *Given a graph G of degree at most d, a vertex cover S of G, and a valid edit sequence α starting from S, we denote by $\mathcal{S}(\alpha) = \{\beta_1, \ldots, \beta_t\}$ the ordered set of d-add-remove segments in α. We say α is a* nice edit sequence *if it is valid, d-well-formed, and satisfies the following invariants:*

- **Connectivity invariant:** *$G[V(\beta_i)]$ is connected for all i between 1 and t.*
- **Early removal invariant:** *For $1 \leq p_1 < p_2 < p_3 \leq |\alpha|$, if $\alpha[p_1] \in \mathcal{E}_a$, $\alpha[p_2] \in \mathcal{E}_a$, and $\alpha[p_3] \in \mathcal{E}_r$, then $V(\alpha[p_3])$ and $V(\alpha[p_1 + 1, p_3 - 1])$ are not separated.*

Intuitively, the early removal invariant states that every removal marker in a nice edit sequence must occur "as early as possible". In other words, a vertex is removed right after its neighbors (and possibly itself) are added. The next observation captures this notion and will be useful to prove the correctness of the ConvertToNice algorithm. We say a removal marker at position $p > 1$ in α satisfies the *non-separation* condition if there exists a position q between 1 and $p - 1$ such that

$\alpha[q] \in \mathcal{E}_a$ and $\alpha[p]$ cannot occur before $\alpha[q]$. Note that if some removal marker at position $p > 1$ in α violates the non-separation condition, then either $\alpha[p]$ belongs to the starting piece of α or α is not nice; as one of the invariants of Definition 3 cannot be satisfied.

Observation 6 *Given a graph G, a vertex cover S of G, and a valid edit sequence α, we let $\alpha[p] = r_i$, where $p > 1$ and $r_i \in \mathcal{E}_r$. If $\{v_i\}$ and $V(\alpha[1, p-1])$ are not separated, then $\alpha[p]$ satisfies the non-separation condition.*

Proof. Since $\{v_i\}$ and $V(\alpha[1, p-1])$ are not separated, one of the following conditions must hold:

(1) $a_i \in \alpha[1, p-1]$,
(2) $r_i \in \alpha[1, p-1]$,
(3) $\{a_j \mid v_j \in N_G(v_i)\} \cap \alpha[1, p-1] \neq \emptyset$, or
(4) $\{r_j \mid v_j \in N_G(v_i)\} \cap \alpha[1, p-1] \neq \emptyset$.

We let $q < p$ denote the closest position to p such that $\alpha[q]$ satisfies one of the conditions above. If $\alpha[q]$ satisfies conditions (1) or (3), we are done. If $\alpha[q]$ satisfies conditions (2) or (4), we have a contradiction since α is valid and a vertex cannot be removed from a vertex cover before itself and all of its neighbors are added. □

The ConvertToNice algorithm can be divided into three stages. In the first stage (Algorithm 1, lines 1–5), the starting piece of α' is updated by copying all possible vertex removals in α that do not satisfy the non-separation condition, i.e. the corresponding removed vertex is separated from all the vertices that are touched in earlier positions in α. Each copied marker is replaced with a $\varnothing$ in α, which is then cut by applying the clean operation. Note that after line 5, $G[V(\alpha')]$ is an independent set in $G[S]$ with no neighbors in $V(G) \setminus S$ as otherwise the vertices in $N_G(V(\alpha')) \cap \{V(G) \setminus S\}$ would have to be added prior to any vertex removal. Hence $V(S, \alpha'[1, p-1])$ is a vertex cover of G. If no other vertices are removed in $clean(\alpha)$ then we are done, since all the vertex additions can be copied to the ending piece of α', which will have an empty central piece. This is handled in the third stage of the algorithm (lines 25–27). Otherwise, we know that each remaining vertex removal in α must satisfy the non-separation condition and is therefore preceded by at least one vertex addition.

The second stage of the algorithm (lines 6–24) handles the central piece of α'. The while loop is repeated as long as $clean(\alpha)$ contains at least one vertex removal. Each iteration handles the first occurrence of a vertex removal in α and copies a d-add-remove segment from α to α'. Combining Observation 6 with the fact that G has degree at most d, we know that at least one and at most d vertex additions have to occur before the corresponding removal can occur. Therefore, the set M_a will contain at most d ordered pairs, where each pair corresponds to a labeled addition marker and its corresponding position in α (lines 7–9). Lines 10–22 will iterate through the subsets of M_a in increasing order of size. For each $N_a \subseteq M_a$, we replace the positions of all addition markers in N_a by $\varnothing$ in α (lines 11–12). If we can find a labeled removal marker which no longer satisfies the non-separation condition after line 12, the addition markers in N_a are copied to α'. Otherwise, the changes done to α are "reversed" (lines 21–22) and we proceed to the next subset of M_a. Whenever the conditions on line 13 are true, lines 16–18 guarantee that each remaining vertex removal in α must again satisfy the non-separation condition. Finally, the third stage of the algorithm (lines 25–27) completes the ending piece of α' by simply copying all remaining vertex additions from α. As an example, consider a graph $G = (V, E)$ where $\{v_1, \ldots, v_6\} \subseteq V(G)$ and

Algorithm 1 ConvertToNice

Input: A graph G of degree at most d, two vertex covers S and T of G, and a valid edit sequence α from S to T.
Output: A nice edit sequence α' from S to T.

1: $p = 1$; $\beta = \alpha$;
2: **for** $(1 \leq i \leq |\alpha|)$
3: **if** $(\alpha[i] = r_x$, for $1 \leq x \leq n$, and $\{v_x\}$ is separated from $V(\beta[1, i-1]))$
4: $\alpha'[p] = \alpha[i]$; $\alpha[i] = \varnothing$; $p = p + 1$;
5: $\alpha = clean(\alpha)$;
6: **while** (α contains at least one element from $\mathcal{E}_r$)
7: **for** $(1 \leq i \leq |\alpha|)$
8: **if** $(\alpha[i] = r_w$, for $1 \leq w \leq n)$
9: $M_a = \{(a_x, q) \mid (v_x \in N_G(v_w) \vee w = x) \wedge a_x \in \alpha[1, i-1] \wedge \alpha[q] = a_x\}$;
10: **for each** $N_a \subseteq M_a$ (in increasing size order)
11: **for each** $(a_x, q) \in N_a$, $1 \leq x \leq n$ and $1 \leq q \leq i-1$
12: $\alpha[q] = \varnothing$;
13: **if** $(\exists j$ s.t. $\alpha[j] = r_y$, $1 \leq y \leq n$, and $\{v_y\}$ is separated from $V(\alpha[1, j-1]))$
14: **for each** $(a_x, q) \in N_a$
15: $\alpha'[p] = a_x$; $p = p + 1$;
16: **for** $(1 \leq l \leq |\alpha|)$
17: **if** $(\alpha[l] = r_z$, $1 \leq z \leq n$, and $\{v_z\}$ is separated from $V(\alpha[1, l-1]))$
18: $\alpha'[p] = \alpha[l]$; $\alpha[l] = \varnothing$; $p = p + 1$;
19: **break**;
20: **else**
21: **for each** $(a_x, q) \in N_a$
22: $\alpha[q] = a_x$;
23: $\alpha = clean(\alpha)$;
24: **end while**
25: **for** $(1 \leq i \leq |\alpha|)$
26: $\alpha'[p] = \alpha[i]$; $p = p + 1$;
27: **return** α';

$\{\{5,1\},\{5,2\},\{5,3\},\{5,4\},\{6,1\},\{6,3\}\} \subseteq E(G)$. In addition, assume $\alpha = \{a_1, a_2, a_3, a_4, r_5, r_6\}$. The ConvertToNice algorithm will produce $\alpha' = \{a_1, a_3, r_6, a_2, a_4, r_5\}$.

Lemma 1. *Given a graph G of degree at most d and two vertex covers S and T of G, the* ConvertToNice *algorithm transforms any valid edit sequence α from S to T into a nice edit sequence α' in $\mathcal{O}(n^4|\alpha|^4 2^d)$ time. Moreover, $|V(S, \alpha'[1,p])| \leq |V(S, \alpha[1,p])|$ for all $1 \leq p \leq |\alpha|$.*

Proof. The fact that α' remains a valid edit sequence from S to T follows from the description of the ConvertToNice algorithm. We show that α' is in fact d-well-formed and satisfies both the connectivity and early removal invariants. If the central piece of α' is empty, the first and third stages of the algorithm guarantee that α' will be d-well-formed. When α' has a non-empty central piece, we show that every iteration of the while loop will copy exactly one d-add-remove segment from α to α'.

First we note that lines 2–4 and 16–18 guarantee that every labeled removal marker in α must satisfy the non-separation condition prior to every iteration of the while loop. Hence, the first marker in α must be an addition marker. We let r_w, $1 \leq w \leq n$, denote the first removal marker in α (line 8). Combining Observation 6 with the fact that G has degree at most d, we know that $|M_a| \leq d$. Note that M_a cannot consist of all the vertices in the closed neighborhood of v_w since this would leave some edge in G uncovered. If we replace all addition markers in M_a by blank markers in α, then r_w will satisfy all the conditions on line 13. Thus, in every iteration of the while

loop, lines 14–19 will execute exactly once. We let $N_a \subseteq M_a$ denote the set which results in the conditions on line 13 being true. Since we iterate over the subsets of M_a in increasing order of size, we know that for the first removal marker r_z satisfying the conditions on line 17, v_z is not separated from any of the vertices in $\{v_x \mid \{a_x, q\} \in N_a\}$. Thus, for every $v_x \in \{v_x \mid \{a_x, q\} \in N_a\}$, $x \neq z$, $\{v_x, v_z\} \in E(G)$. Moreover, if r_z is at position q_1, $1 \leq q_1 \leq |\alpha|$, any other removal marker at position $q_2 > q_1$, $q_2 \leq |\alpha|$, which now satisfies the conditions on line 17 is also not separated from any of the vertices in $\{v_x \mid \{a_x, q\} \in N_a\}$; otherwise $|N_G[v_z] \cap M_A| < |N_a|$, contradicting the fact that N_a is the smallest subset of M_a which results in the conditions on line 13 being true. Finally, we note that since we iterate over the subsets of M_a in increasing order of size, at most d removal markers will satisfy the conditions on line 17. Putting it all together, lines 14–19 will execute exactly once and copy at most d addition markers followed by at most d removal markers. We let β denote this d-add-remove segment.

The fact that $G[V(\beta)]$ is connected follows from the observation that every removed vertex in β is not separated from any of the added vertices in β. The early removal invariant is satisfied because removal markers are copied from α to α' as soon as they no longer satisfy the non-separation condition (lines 2–4 and 16–18). We conclude that α' is a nice edit sequence from S to T.

The $\mathcal{O}(n^4|\alpha|^4 2^d)$ bound on the running time follows from the description of the ConvertToNice algorithm. To prove $|V(S, \alpha'[1,p])| \leq |V(S, \alpha[1,p])|$, for all $1 \leq p \leq |\alpha|$, we show that the number of removal markers in $\alpha'[1,p]$ is greater than or equal to the number of removal markers in $\alpha[1,p]$. We assume without loss of generality that $\alpha[1] \in \mathcal{E}_a$, as all removal markers which are not preceded by an addition marker in α will be copied to α' in the first stage of the algorithm. In lines 13–19 of the algorithm, the number of addition markers copied from α to α' is equal to $|N_a| \leq |M_a|$. Moreover, whenever any marker is shifted from position q_1 in α to position $q_2 > q_1$ in α', it must be the case that some removal marker was shifted from position $q_3 > q_1$ in α to some position $q_4 < q_2$ in α'. In other words, if any marker is "shifted to the right" when copied from α to α' it must be the case that a removal marker was "shifted" to the left. Hence, there exists no p such that $\alpha[1,p]$ contains more removal markers than $\alpha'[1,p]$. □

3 Hardness Results

In earlier work establishing the $\mathbf{W}[\mathbf{1}]$-hardness of the VCR problem parameterized by ℓ on general graphs, it was also shown that the problem becomes fixed-parameter tractable whenever $\ell = |S_R \cup T_A|$ [21]. When $|S_R \cup T_A| = n$, we know from Observation 2 that $\ell \geq n$, since every vertex in $S_R \cup T_A$ must be touched at least once. Therefore, even though the problem is fixed-parameter tractable, any algorithm solving an instance of the VCR problem where $|S_R \cup T_A| = \ell = n$ would run in time exponential in n. Moreover, Observation 1 implies that whenever $|S_R \cup T_A| = n$ the input graph must be bipartite. It is thus natural to ask about the complexity of the problem when $|S_R \cup T_A| = \ell < n$ and the input graph is restricted to be bipartite. Since the Vertex Cover problem is known to be solvable in time polynomial in n on bipartite graphs, our result is, to the best of our knowledge, the first example of a problem solvable in polynomial time whose reconfiguration version is $\mathbf{W}[\mathbf{1}]$-hard.

For a graph $G = (V, E)$, a *crown* is a pair (W, H) satisfying the following properties: (i) $W \neq \emptyset$ is an independent set of G, (ii) $N_G(W) = H$, and (iii) there exists a matching in $G[W \cup H]$ which saturates H [1, 4]. H is called the *head* of the crown and the *width* of the crown is $|H|$. Crown structures have played a central role in the development of efficient kernelization algorithms for the Vertex Cover problem [1, 4]. We define the closely related notion of (k, d)-constrained crowns and

show in the remainder of this section that the complexity of finding such structures in a bipartite graph is central for determining the complexity of the reconfiguration problem.

For a graph $G = (V, E)$, we define a (k, d)-*constrained crown* as a pair (W, H) satisfying all properties of a regular crown with the additional constraints that $|H| \leq k$ and $|W| - |H| \geq d \geq 0$. We are now ready to introduce the (k, d)-Bipartite Constrained Crown problem, or (k, d)-BCC, which is formally defined as follows:

(k, d)-Bipartite Constrained Crown
Input: A bipartite graph $G = (A \cup B, E)$ and two positive integers k and d
Parameters: k and d
Question: Does G have a (k, d)-constrained crown (W, H) such that $W \subseteq A$ and $H \subseteq B$?

Lemma 2. (k, d)-Bipartite Constrained Crown *is* $\mathbf{W}[\mathbf{1}]$*-hard even when the input graph,* $G = (A \cup B, E)$*, is* C_4*-free and all vertices in* A *have degree at most two.*

Proof. We give an **FPT** reduction from k-Clique, known to be $\mathbf{W}[\mathbf{1}]$-hard, to $(k, \binom{k}{2})$-Bipartite Constrained Crown. For (G, k) an instance of k-Clique, we let $V(G) = \{v_1, \ldots, v_n\}$ and $E(G) = \{e_1, \ldots, e_m\}$.

We first form a bipartite graph $G' = ((X \cup Z) \cup Y, E_1 \cup E_2)$, where vertex sets X and Y contain one vertex for each vertex in $V(G)$ and Z contains one vertex for each edge in $E(G)$. More formally, we set $X = \{x_1, \ldots, x_n\}$, $Y = \{y_1, \ldots, y_n\}$, and $Z = \{z_1, \ldots, z_m\}$. The edges in E_1 join each pair of vertices x_i and y_i for $1 \leq i \leq n$ and the edges in E_2 join each vertex z in Z to the two vertices y_i and y_j corresponding to the endpoints of the edge in $E(G)$ to which z corresponds. Since each edge either joins vertices in X and Y or vertices in Y and Z, it is not difficult to see that the vertex sets $X \cup Z$ and Y form a bipartition.

By our construction, G' is C_4-free; vertices in X have degree 1, and since there are no double edges in G, i.e. two edges between the same pair of vertices, no pair of vertices in Y can have more than one common neighbour in Z.

For $(G', k, \binom{k}{2})$ an instance of $(k, \binom{k}{2})$-BCC, $A = X \cup Z$, and $B = Y$, we claim that G has a clique of size k if and only if G' has a $(k, \binom{k}{2})$-constrained crown (W, H) such that $W \subseteq A$ and $H \subseteq B$.

If G has a clique K of size k, we set $H = \{y_i \mid v_i \in V(K)\}$, namely the vertices in Y corresponding to the vertices in the clique. To form W, we choose $\{x_i \mid v_i \in V(K)\} \cup \{z_i \mid e_i \in E(K)\}$, that is, the vertices in X corresponding to the vertices in the clique and the vertices in Z corresponding to the edges in the clique. Clearly H is a subset of size k of B and W is a subset of size $k + \binom{k}{2}$ of A; this implies that $|W| - |H| \geq d = \binom{k}{2}$, as required. To see why $N_{G'}(W) = H$, it suffices to note that every vertex $x_i \in W$ is connected to exactly one vertex $y_i \in H$ and every degree-two vertex $z_i \in W$ corresponds to an edge in K whose endpoints $\{v_i, v_j\}$ must have corresponding vertices in H. Moreover, due to E_1 there is a matching between the vertices of H and the vertices of W in X, and hence a matching in $G'[W \cup H]$ which saturates H.

We now assume that G' has a $(k, \binom{k}{2})$-constrained crown (W, H) such that $W \subseteq X \cup Z$ and $H \subseteq Y$. It suffices to show that $|H|$ must be equal to k, $|W \cap Z|$ must be equal to $\binom{k}{2}$, and hence $|W \cap X|$ must be equal to k; from this we can conclude the vertices in $\{v_i \mid y_i \in H\}$ form a clique of size k in G as $|W \cap Z| = \binom{k}{2}$, requiring that edges exist between each pair of vertices in the set

$\{v_i \mid y_i \in H\}$. Moreover, since $|W \cap X| = k$ and $N_{G'}(W) = H$, a matching that saturates H can be easily found by simply picking all edges $\{x_i, y_i\}$ for $y_i \in H$.

To prove the sizes of H and W, we first observe that since $|H| \leq k$, $N_{G'}(W) = H$, and each vertex in Y has exactly one neighbour in X, we know that $|W \cap X| \leq |H| \leq k$. Moreover, since $|W| = |W \cap X| + |W \cap Z|$ and $|W| - |H| \geq \binom{k}{2}$, we know that $|W \cap Z| = |W| - |W \cap X| \geq \binom{k}{2} + |H| - |W \cap X| \geq \binom{k}{2}$. If $|W \cap Z| = \binom{k}{2}$ our proof is complete since, by our construction of G', H is a set of at most k vertices in the original graph G and the subgraph induced by those vertices in G has $\binom{k}{2}$ edges. Hence, $|H|$ must be equal to k. Suppose instead that $|W \cap Z| > \binom{k}{2}$. In this case, since each vertex of Z has degree two, the number of neighbours of $W \cap Z$ in Y is greater than k, violating the assumptions that $N_{G'}(W) = H$ and $|H| \leq k$. □

We can now show the main result of this section:

Theorem 1. Vertex Cover Reconfiguration *parameterized by ℓ and restricted to bipartite graphs is* $\mathbf{W}[\mathbf{1}]$*-hard .*

Proof. We give an **FPT** reduction from (t, d)-Bipartite Constrained Crown to Vertex Cover Reconfiguration in bipartite graphs. For $(G = (A \cup B, E), t, d)$ an instance of (t, d)-Bipartite Constrained Crown, $A = \{a_1, \ldots, a_{|A|}\}$, and $B = \{b_1, \ldots, b_{|B|}\}$, we form $G' = (X \cup Y \cup U \cup V, E_1 \cup E_2)$ such that X and Y correspond to the vertex sets A and B, E_1 connects vertices in X and Y corresponding to vertices in A and B joined by edges in G, and U, V, and E_2 form a bipartite clique $K_{d+t,d+t}$. More formally, $X = \{x_1, \ldots, x_{|A|}\}$, $Y = \{y_1, \ldots, y_{|B|}\}$, $U = \{u_1, \ldots, u_{d+t}\}$, $V = \{v_1, \ldots, v_{d+t}\}$, $E_1 = \{\{x_i, y_j\} \mid \{a_i, b_j\} \in E(G)\}$ and $E_2 = \{\{u_i, v_j\} \mid 1 \leq i \leq d + t, 1 \leq j \leq d + t\}$.

We let $(G', S, T, k = |A| + d + 2t, \ell = 4d + 6t)$ be an instance of Vertex Cover Reconfiguration, where $S = X \cup U$ and $T = X \cup V$. Clearly $|S| = |T| = |A| + d + t$. We claim that G has a (k, d)-constrained crown (W, H) such that $W \subseteq A$ and $H \subseteq B$ if and only if there is a path of length less than or equal to $4d + 6t$ from S to T.

If G has such a pair (W, H), we form a reconfiguration sequence of length at most $4d + 6t$ as follows:

1. Add each vertex y_i such that $b_i \in H$. The resulting vertex cover size is $|A| + d + t + |H|$.
2. Remove $d + |H|$ vertices x_i such that $a_i \in W$. The resulting vertex cover size is $|A| + t$.
3. Add each vertex from V. The resulting vertex cover size is $|A| + d + 2t$.
4. Remove each vertex from U. The resulting vertex cover size is $|A| + t$.
5. Add each vertex removed in phase 2. The resulting vertex cover size is $|A| + d + t + |H|$.
6. Remove each vertex added in phase 1. The resulting vertex cover size is $|A| + d + t$.

The length of the sequence follows from the fact that $|H| \leq t$: phases 1 and 6 consist of at most t steps each and phases 2, 3, 4, and 5 of at most $d + t$ steps each. The fact that each set forms a vertex cover is a consequence of the fact that $N_G(W) = H$.

For the converse, we observe that before removing any vertex u_i, $1 \leq i \leq d + t$, from U, we first need to add all $d + t$ vertices from V. Therefore, if there is a path of length at most $4d + 6t$ from S to T, then we can assume without loss of generality that there exists a node Q (i.e. a vertex cover) along this path such that:

(1) $|Q| \leq |A| + t$ and,
(2) all vertices that were touched in order to reach node Q belong to $X \cup Y$.

In other words, at node Q, the available capacity is greater than or equal to $d+t$ and all edges in $G[U \cup V]$ are still covered by U. We let $Q_{\text{IN}} = Q \setminus S$ and $Q_{\text{OUT}} = S \setminus Q$. Since $S = X \cup U$, $Q_{\text{IN}} \subseteq Y$ and $Q_{\text{OUT}} \subseteq X$. Moreover, Since $|Q| = |S| + |Q_{\text{IN}}| - |Q_{\text{OUT}}| = |A| + d + t + |Q_{\text{IN}}| - |Q_{\text{OUT}}| \leq |A| + t$, we know that $|Q_{\text{OUT}}| - |Q_{\text{IN}}|$ must be greater than or equal to d. Given that $\ell \leq 4d + 6t$ and we need exactly $2d + 2t$ steps to add all vertices in V and remove all vertices in U, we have $2d + 4t$ remaining steps to allocate elsewhere. Therefore, $|Q_{\text{OUT}}| + |Q_{\text{IN}}| \leq d + 2t$ as $Q_{\text{IN}} \subseteq Y$, $Q_{\text{OUT}} \subseteq X$, and every vertex in $Q_{\text{IN}} \cup Q_{\text{OUT}}$ must be touched at least twice (i.e. added and then removed). Combining those observations, we get:

$|Q_{\text{OUT}}| + |Q_{\text{IN}}| \leq d + 2t$
$|Q_{\text{IN}}| - |Q_{\text{OUT}}| \leq -d$
$|Q_{\text{IN}}| \leq t$

We have just shown that G has a pair $(Q_{\text{OUT}}, Q_{\text{IN}})$ such that $Q_{\text{OUT}} \subseteq X$, $Q_{\text{IN}} \subseteq Y$, $|Q_{\text{IN}}| \leq t$, $|Q_{\text{OUT}}| - |Q_{\text{IN}}| \geq d \geq 0$, and $N_G(Q_{\text{OUT}}) = Q_{\text{IN}}$ as otherwise some edge is not covered. The remaining condition for $(Q_{\text{OUT}}, Q_{\text{IN}})$ to satisfy is for $G[Q_{\text{OUT}} \cup Q_{\text{IN}}]$ to have a matching which saturates Q_{IN}. Hall's Marriage Theorem [13] states that such a saturating matching exists if and only if for every subset P of Q_{IN}, $|P| \leq |N_{G[Q_{\text{OUT}} \cup Q_{\text{IN}}]}(P)|$. By a simple application of Hall's theorem, if no such matching exists then there exists a subgraph Z of $G[Q_{\text{OUT}} \cup Q_{\text{IN}}]$ such that $|V(Z) \cap Q_{\text{OUT}}| < |V(Z) \cap Q_{\text{IN}}|$. By deleting this subgraph from $Q_{\text{OUT}} \cup Q_{\text{IN}}$ we can get a new pair $(Q'_{\text{OUT}}, Q'_{\text{IN}})$ which must still satisfy $Q'_{\text{OUT}} \subseteq X$, $Q'_{\text{IN}} \subseteq Y$, $|Q'_{\text{IN}}| \leq t$, $|Q'_{\text{OUT}}| - |Q'_{\text{IN}}| \geq d \geq 0$, and $N_G(Q'_{\text{OUT}}) = Q'_{\text{IN}}$ since we delete more vertices from Q_{IN} than we do from Q_{OUT} and $N_{G[Q_{\text{OUT}} \cup Q_{\text{IN}}]}(V(Z) \cap Q_{\text{IN}}) = V(Z) \cap Q_{\text{OUT}}$. Finally, if $(Q'_{\text{OUT}}, Q'_{\text{IN}})$ does not have a matching which saturates Q'_{IN}, we can repeatedly apply the same rule until we reach a pair which satisfies all the required properties. Since $|Q_{\text{OUT}}| \geq |Q_{\text{IN}}|$, such a pair is guaranteed to exist as otherwise every subset P of Q_{IN} would satisfy $|P| > |N_{G[Q_{\text{OUT}} \cup Q_{\text{IN}}]}(P)|$ and hence $|Q_{\text{OUT}}| < |Q_{\text{IN}}|$, a contradiction. □

4 Polynomial-Time Algorithms

In this section, we present a characterization of instances of the VCR problem solvable in time polynomial in n, and apply this characterization to trees and cactus graphs. In both cases, we show how to find reconfiguration sequences of shortest possible length and therefore ignore the parameter ℓ. Unless stated otherwise, reconfiguration sequences are represented as ordered sequences of vertex covers or nodes in the reconfiguration graph.

Definition 4. *Given two vertex covers of G, A and B, a reconfiguration sequence β from A to some vertex cover A' is a* c-bounded prefix *of a reconfiguration sequence α from A to B, if and only if all of the following conditions hold:*

(1) $|A'| \leq |A|$.
(2) For every node A'' in β, $|A''| \leq |A| + c$.
(3) For every node A'' in β, A'' is obtained from its predecessor by either the removal or the addition of a single vertex in the symmetric difference of the predecessor and B.
(4) No vertex is touched more than once in the course of β.

We write $A \overset{c,B}{\longleftrightarrow} A'$ when such a c-bounded prefix exists.

Observation 7 *Given two vertex covers S and T of G, if G has a vertex cover S' such that $S \overset{c,T}{\longleftrightarrow} S'$, then $S \overset{d,T}{\longleftrightarrow} S'$ for all $d > c$.*

Lemma 3. *Given two vertex covers S and T of G and two positive integers k and c such that $|S|, |T| \leq k$, a reconfiguration sequence α of length $|S_R| + |T_A| = |S \Delta T|$ from S to T exists if:*

(1) $|S| \leq k - c$,
(2) $|T| \leq k - c$, and
(3) For any two vertex covers A and B of G such that $|A| \leq k-c$ and $|B| \leq k-c$, either $A \overset{c,B}{\longleftrightarrow} A'$ or $B \overset{c,A}{\longleftrightarrow} B'$, where A' and B' are vertex covers of G.

Moreover, if c-bounded prefixes can be found in time polynomial in n, then α can be found in time polynomial in n.

Proof. We prove the lemma by induction on $|S \Delta T|$. When $|S \Delta T| = 0$, S is equal to T, and the claim holds trivially since $|\alpha| = 0$.

When $|S \Delta T| > 0$, we know that either $S \overset{c,T}{\longleftrightarrow} S'$ or $T \overset{c,S}{\longleftrightarrow} T'$. Without loss of generality, we assume $S \overset{c,T}{\longleftrightarrow} S'$ and let β denote the c-bounded prefix from S to S'. From Definition 4, we know that the size of every node in β is no greater than $|S| + c \leq k$. Therefore, the maximum allowed capacity constraint is never violated.

Since $|S'| \leq |S|$ (Definition 4), by the induction hypothesis there exists a reconfiguration sequence from S' to T whose length is $|S' \Delta T|$. By appending the reconfiguration sequence from S' to T to the reconfiguration sequence from S to S', we obtain a reconfiguration sequence α from S to T.

To show that $|\alpha| = |S \Delta T|$, it suffices to show that $|\beta| + |S' \Delta T| = |S \Delta T|$. We know that no vertex is touched more than once in β and every touched vertex belongs to $S \Delta T$ (Definition 4). We let $H \subseteq S \Delta T$ denote the set of touched vertices in β and we subdivide H into two sets $H_S = H \cap S = H \cap S_R$ and $H_T = H \cap T = H \cap T_A$. It follows that $|\beta| = |H_S| + |H_T|$ and $|S' \Delta T| = |S_R \setminus H_S| + |T_A \setminus H_T|$. Therefore, $|\beta| + |S' \Delta T| = |H_S| + |H_T| + |S_R \setminus H_S| + |T_A \setminus H_T| = |S_R| + |T_A| = |S \Delta T|$ as needed.

When c-bounded prefixes can be found in time polynomial in n, the proof gives an algorithm for constructing the full reconfiguration sequence from S to T in time polynomial in n. □

4.1 Trees

Theorem 2. Vertex Cover Reconfiguration *restricted to trees can be solved in time polynomial in n.*

Proof. We let (G, S, T, k, ℓ) be an instance of Vertex Cover Reconfiguration. The proof proceeds in two stages. We start by showing that when G is a tree and S and T are of size at most $k - 1$, we can always find 1-bounded prefixes $S \overset{1,T}{\longleftrightarrow} S'$ or $T \overset{1,S}{\longleftrightarrow} T'$ in time polynomial in n. Therefore, we can apply Lemma 3 with $c = 1$ to find a reconfiguration sequence of length $|S \Delta T|$ from S to T in time polynomial in n. In the second part of the proof, we show how to handle the remaining cases where S, T, or both S and T are of size greater than $k - 1$.

First, we note that every forest either has a degree-zero or a degree-one vertex. Hence, trees and forests are 1-degenerate graphs. Since G is a tree, $G[S_R \cup T_A]$ is a forest and is therefore 1-degenerate. To find 1-bounded prefixes in $G[S_R \cup T_A]$, it is enough to find a vertex of degree at most one, which can clearly be done in time polynomial in n: For any two vertex covers S and T of a tree G such that $S, T \leq k-1$, we can always find a vertex $v \in S_R \cup T_A$ having degree at most one in $G[S_R \cup T_A]$. The existence of v guarantees the existence of a 1-bounded prefix from either S to some vertex cover S' or from T to some vertex cover T'. When $v \in S_R$ and $|N_{G[S_R \cup T_A]}(v)| = 0$, we have $S \xleftrightarrow{0,T} S'$ since S' is obtained from S by simply removing v. When $v \in S_R$ and $|N_{G[S_R \cup T_A]}(v)| = 1$, we have $S \xleftrightarrow{1,T} S'$ since S' is obtained from S by first adding the unique neighbor of v and then removing v. Similar arguments hold when $v \in T_A$.

Therefore, combining Lemma 3 and the fact that $G[S_R \cup T_A]$ is 1-degenerate, we know that if $|S| \leq k-1$ and $|T| \leq k-1$, a reconfiguration sequence of length $|S_R| + |T_A|$ from S to T can be found in time polynomial in n. Furthermore, since the length of a reconfiguration sequence can never be less than $|S_R| + |T_A|$, the reconfiguration sequence given by Lemma 3 is a shortest path from S to T in the reconfiguration graph.

When S (or T) has size k and is minimal, then we have a no-instance since neither removing nor adding a vertex results in a k-vertex cover, and hence S (or T) will be an isolated node in the reconfiguration graph, with no path to T (or S).

When S, T, or both S and T are of size k and are non-minimal, there always exists a reconfiguration sequence from S to T, since S and T can be reconfigured to solutions S' and T', respectively, of size less than k, to which Lemma 3 can be applied. The only reconfiguration steps from S (or T) of size k are to subsets of S of size $k-1$ (or to subsets of T of size $k-1$); the reconfiguration sequence obtained from Lemma 3 is thus a shortest path. Therefore, we can obtain a shortest path from S to T through a careful selection of S' and T'. There are two cases to consider:

Case (1): $|S| = k$, $|T| = k$, S is non-minimal, and T is non-minimal. When both S and T are of size k and are non-minimal, then each must contain at least one removable vertex. Hence, by removing such vertices, we can transform S and T into vertex covers S' and T', respectively, of size $k-1$. We let u and v be removable vertices in S and T respectively, and we set $S' = S \setminus \{u\}$ and $T' = T \setminus \{v\}$.

1. If $u \in S_R$ and $v \in T_A$, then the length of a shortest reconfiguration sequence from S' to T' will be $|S' \Delta T'| = |S \Delta T| - 2$. Therefore, accounting for the two additional removals, the length of a shortest path from S to T will be equal to $|S \Delta T|$.
2. If $u \in S_R$ and $v \in C_{ST}$, then the length of a shortest reconfiguration sequence from S' to T' will be $|S' \Delta T'| = |S \Delta T| - 1$. Since v is in C_{ST}, it must be removed and added back. Therefore, the length of a shortest path from S to T will be equal to $|S \Delta T| + 2$. The same is true when $u \in C_{ST}$ and $v \in T_A$ or when $u = v$ and $u \in C_{ST}$.
3. Otherwise, when $u \in C_{ST}$, $v \in C_{ST}$, and $u \neq v$, the length of a shortest path from S to T will be $|S \Delta T| + 4$ since we have to touch two vertices in C_{ST} (i.e. two extra additions and two extra removals).

Case (2): $|S| = k$, $|T| = k-1$, and S is non-minimal. Since $|T| = k-1$, we only need to

reduce the size of S to $k-1$ in order to apply Lemma 3. Since S is non-minimal, it must contain at least one removable vertex. We let u be a removable vertex in S and we set $S' = S \setminus \{u\}$.

1. If $u \in S_R$, then the length of a shortest reconfiguration sequence from S' to T will be $|S'\Delta T| = |S\Delta T| - 1$. Therefore, accounting for the additional removal, the length of a shortest path from S to T will be equal to $|S\Delta T|$.
2. If $u \in C_{ST}$, then the length of a shortest reconfiguration sequence from S' to T will be $|S'\Delta T| = |S\Delta T|$. Since v is in C_{ST}, it must be removed and added back. Therefore, the length of a shortest path from S to T will be equal to $|S\Delta T| + 2$.

Similar arguments hold for the symmetric case where $|S| = k-1$, $|T| = k$, and T is non-minimal.

As there are at most k^2 pairs of removable vertices in S and T to check for Case (1), we can exhaustively try all pairs and choose one that minimizes the length of a reconfiguration sequence. Similarly, there are at most k removable vertices to check in Case (2). Consequently, Vertex Cover Reconfiguration in trees can be solved in time polynomial in n. □

4.2 Cactus Graphs

A *cactus graph* G is a connected graph in which every edge belongs to at most one cycle. We let $\mathcal{C}(G)$ denote the set of all cycles in G. We say vertex $v \in V(G)$ is a *join vertex* if v belongs to a cycle and $N_G(v) \geq 3$. The following observation is a consequence of the fact that any maximal matching $\mathcal{M}(G)$ in a cactus graph G can contain an edge from each cycle in $\mathcal{C}(G)$.

Observation 8 *For a cactus graph G, the number of cycles in G is bounded above by the size of any maximal matching $\mathcal{M}(G)$, i.e. $|\mathcal{C}(G)| \leq |\mathcal{M}(G)|$.*

The next observation is a consequence of the fact that for any cactus graph G, we can obtain a spanning tree of G by removing a single edge from every cycle in G.

Observation 9 *For a cactus graph G and T_G a spanning tree of G, the total number of edges in G is equal to the number of edges in T_G plus the total number of cycles in G, i.e. $|E(G)| = |E(T_G)| + |\mathcal{C}(G)| = |V(T_G)| - 1 + |\mathcal{C}(G)|$.*

Any graph with no even cycles is a cactus graph. For a graph G with no even cycles and any two vertex covers, S and T, of G, we know that $G[S_R \cup T_A]$ must be a forest, i.e. a bipartite graph with no even cycles (Observation 1). Observation 10 follows from the fact that in the proof of Theorem 2, the fact that G is a tree is used only to determine that $G[S_R \cup T_A]$ must be a forest. Therefore, using the same proof as in Theorem 2, we can show:

Observation 10 Vertex Cover Reconfiguration *on graphs with no even cycles can be solved in time polynomial in n.*

The goal of this section is to generalize Observation 10 to all cactus graphs. To do so, we first show, in Lemmas 4 and 5, that the third condition of Lemma 3 is satisfied for cactus graphs with $c = 2$. In Lemma 6, we show how 2-bounded prefixes can be found in time polynomial in n, which leads to Theorem 3, the main result of this section.

Lemma 4. *Given two vertex covers S and T of G, there exists a vertex cover S' (or T') of G such that $S \overset{2,T}{\longleftrightarrow} S'$ (or $T \overset{2,S}{\longleftrightarrow} T'$) if one of the following conditions holds:*

(1) $G[S_R \cup T_A]$ has a vertex $v \in S_R$ ($v \in T_A$) such that $|N_{G[S_R \cup T_A]}(v)| \leq 1$, or
(2) there exists a cycle Y in $G[S_R \cup T_A]$ such that all vertices in $Y \cap S_R$ ($Y \cap T_A$) have degree exactly two in $G[S_R \cup T_A]$.

Moreover, both conditions can be checked in time polynomial in n and when one of them is true the corresponding 2-bounded prefix can be found in time polynomial in n.

Proof. First, we note that checking for condition (1) can be accomplished in time polynomial in n by simply inspecting the degree of every vertex in $G[S_R \cup T_A]$. The total number of cycles satisfying condition (2) is linear in the number of degree-two vertices in $G[S_R \cup T_A]$. Therefore, we can check for condition (2) in time polynomial in n by a simple breadth-first search starting from every degree-two vertex in $G[S_R \cup T_A]$.

If $G[S_R \cup T_A]$ has a vertex $v \in S_R$ of degree zero, we let S' denote the vertex cover obtained by simply removing v from S. It is easy to see that the reconfiguration sequence from S to S' is a 0-bounded prefix and can be found in time polynomial in n.

Similarly, if $G[S_R \cup T_A]$ has a vertex $v \in S_R$ of degree one, we let S' denote the node obtained by the addition of the single vertex in $N_{G[S_R \cup T_A]}(v)$ followed by the removal of v. The reconfiguration sequence from S to S' is a 1-bounded prefix and can be found in time polynomial in n.

For the second case, we let Y be a cycle in $G[S_R \cup T_A]$ and we partition the vertices of the cycle into two sets; $Y_S = Y \cap S_R$ and $Y_T = Y \cap T_A$. Since $G[S_R \cup T_A]$ is bipartite, we know that $|Y_S| = |Y_T|$. Since all vertices in Y_S have degree exactly two in $G[S_R \cup T_A]$, it follows that $N_{G[S_R \cup T_A]}(Y_S) \subseteq Y_T$. Therefore, a reconfiguration sequence from S to some vertex cover S' that adds all vertices in Y_T (one by one), and then removes all vertices in Y_S (one by one) will satisfy conditions (1), (3), and (4) from Definition 4 for any value of c. For $c = 2$, such a sequence will not satisfy condition (2) if the cycle has at least six vertices (i.e. $|Y_T| \geq 3$). However, using the fact that every vertex in Y_S has degree exactly two in $G[S_R \cup T_A]$, we can find a reconfiguration sequence from S to S' in which no vertex cover has size greater than $|S| + 2$. To do so, we restrict our attention to $G[Y_S \cup Y_T]$. Since Y is an even cycle, we can label all the vertices of Y in clockwise order from 0 to $|Y| - 1$ such that all vertices in Y_S receive even labels. The reconfiguration sequence from S to S' starts by adding the two vertices labeled 1 and $|Y| - 1$. After doing so, the vertex labeled 0 is removed. Next, to remove the vertex labeled 2, we only need to add the vertex labeled 3. The same process is repeated for all vertices with even labels up to $|Y| - 4$. Finally, when we reach the vertex labeled $|Y| - 2$, both of its neighbors will have already been added and we can simply remove it. Hence, we have a 2-bounded prefix from S to S' and it is not hard to see that finding this reconfiguration sequence can be accomplished in time polynomial in n.

When the appropriate assumptions hold, we can show the symmetric case $T \overset{2,S}{\longleftrightarrow} T'$ using similar arguments. □

Lemma 5. *If G is a cactus graph and S and T are two vertex covers of G, then there exists a vertex cover S' (or T') of G such that $S \overset{2,T}{\longleftrightarrow} S'$ (or $T \overset{2,S}{\longleftrightarrow} T'$).*

Proof. We assume that $|S_R| \geq |T_A|$, as we can swap the roles of S and T whenever $|S_R| < |T_A|$. We observe that every connected component of $G[S_R \cup T_A]$ is a cactus graph since every induced

subgraph of a cactus graph is also a cactus graph. Since we assume $|S_R| \geq |T_A|$, at least one connected component X of $G[S_R \cup T_A]$ must satisfy $|V(X) \cap S_R| \geq |V(X) \cap T_A|$.

To prove the lemma, we show that if neither condition of Lemma 4 applies to X, it must be the case that $|V(X) \cap S_R| < |V(X) \cap T_A|$, contradicting our assumption. To simplify notation, we assume without loss of generality that $G[S_R \cup T_A]$ is connected, as we can otherwise set $G[S_R \cup T_A] = X$. The proof proceeds in two steps. First, we show that if condition (1) of Lemma 4 is not satisfied, then $G[S_R \cup T_A]$ must have at least one vertex $u \in S_R$ of degree at most two in $G[S_R \cup T_A]$. In the second step, we show that if both conditions (1) and (2) of Lemma 4 are not satisfied, then $|S_R| < |T_A|$, which completes the proof by contradiction.

Since $G[S_R \cup T_A]$ is a cactus graph, we can apply Observations 8 and 9 to get:

$$\begin{aligned} |E(G[S_R \cup T_A])| &= |S_R| + |T_A| - 1 + |\mathcal{C}(G[S_R \cup T_A])| \\ &\leq |S_R| + |T_A| - 1 + |\mathcal{M}(G[S_R \cup T_A])| \end{aligned} \tag{1}$$

Moreover, since $G[S_R \cup T_A]$ is bipartite (Observation 1), the size of a maximum matching in $G[S_R \cup T_A]$ is less than or equal to $min(|S_R|, |T_A|)$. Therefore:

$$|\mathcal{C}(G[S_R \cup T_A])| \leq |\mathcal{M}(G[S_R \cup T_A])| \leq |S_R| \tag{2}$$

Combining (1) and (2), we get:

$$\begin{aligned} |E(G[S_R \cup T_A])| &= |S_R| + |T_A| - 1 + \mathcal{C}(G[S_R \cup T_A]) \\ &\leq 2|S_R| + |T_A| - 1 \end{aligned} \tag{3}$$

If the minimum degree in $G[S_R \cup T_A]$ of any vertex in S_R is three or more, then $3|S_R| \leq |E(G[S_R \cup T_A])| \leq 2|S_R| + |T_A| - 1$ and thus $|S_R| \leq |T_A| - 1$, contradicting our assumption that $|S_R| \geq |T_A|$. Hence, $G[S_R \cup T_A]$ must have at least one vertex of degree two in S_R.

Next, we show that if $G[S_R \cup T_A]$ has no vertex $v \in S_R$ such that $|N_{G[S_R \cup T_A]}(v)| \leq 1$ and no cycle Y such that all vertices in $Y \cap S_R$ have degree exactly two in $G[S_R \cup T_A]$, then $|S_R| < |T_A|$. We let S^x denote the set of vertices in S_R having degree x in $G[S_R \cup T_A]$. Since $G[S_R \cup T_A]$ has no vertex $v \in S_R$ such that $|N_{G[S_R \cup T_A]}(v)| \leq 1$, we know that S^2 cannot be empty. In addition, since there is no cycle Y in $G[S_R \cup T_A]$ such that all vertices in $Y \cap S_R$ have degree exactly two in $G[S_R \cup T_A]$, any cycle involving a vertex in S^2 must also include a vertex from $\bigcup_{i \geq 3} S^i$. It follows that $\bigcup_{i \geq 3} S^i$ is a feedback vertex set of $G[S_R \cup T_A]$ and $G[S^2 \cup T_A]$ is a forest.

We let m_s denote the maximum degree in $G[S_R \cup T_A]$ of any vertex in S_R. Since each edge in $G[S_R \cup T_A]$ has one endpoint in S_R,

$$\sum_{i=2}^{m_s} i|S^i| \leq |E(G[S_R \cup T_A])| \tag{4}$$

and since each vertex in S_R is in some S^i, and using (1), we can rewrite (4) as

$$\sum_{i=2}^{m_s} i|S^i| \leq \left(\sum_{i=2}^{m_s} |S^i| \right) + |T_A| - 1 + |\mathcal{C}(G[S_R \cup T_A])|. \tag{5}$$

To bound $|\mathcal{C}(G[S_R \cup T_A])|$, we note that since no edge can belong to more than one cycle in a cactus graph, any vertex $v \in S^x$ can be involved in at most $\lfloor \frac{x}{2} \rfloor$ cycles. Combining this observation with the fact that any cycle involving a vertex in S^2 must also include a vertex from $\bigcup_{i \geq 3} S^i$, we have:

$$\begin{aligned}\sum_{i=2}^{m_s} i|S^i| &\leq \left(\sum_{i=2}^{m_s} |S^i|\right) + |T_A| - 1 + \left(\sum_{i=3}^{m_s} \lfloor \frac{i}{2} \rfloor |S^i|\right) \\ &\leq |S^2| + \left(\sum_{i=3}^{m_s} (1 + \lfloor \frac{i}{2} \rfloor)|S^i|\right) + |T_A| - 1 \end{aligned} \tag{6}$$

Finally, by rewriting $\sum_{i=2}^{m_s} i|S^i|$ as $2|S^2| + \sum_{i=3}^{m_s} i|S^i|$ and given that $i - (1 + \lfloor \frac{i}{2} \rfloor) \geq 1$ for $i \geq 3$, we obtain the desired bound:

$$\begin{aligned} 2|S^2| + \sum_{i=3}^{m_s} i|S^i| &\leq |S^2| + \left(\sum_{i=3}^{m_s} (1 + \lfloor \frac{i}{2} \rfloor)|S^i|\right) + |T_A| - 1 \\ |S^2| + \sum_{i=3}^{m_s} i|S^i| &\leq \left(\sum_{i=3}^{m_s} (1 + \lfloor \frac{i}{2} \rfloor)|S^i|\right) + |T_A| - 1 \\ |S^2| + \sum_{i=3}^{m_s} (i - (1 + \lfloor \frac{i}{2} \rfloor))|S^i| &\leq |T_A| - 1 \\ |S_R| = \sum_{i=2}^{m_s} |S^i| &\leq |T_A| - 1 \end{aligned} \tag{7}$$

This completes the proof. □

Lemma 6. *If G is a cactus graph and S and T are vertex covers of G, then finding a 2-bounded prefix from S to some vertex cover S' (or from T to some vertex cover T') of G can be accomplished in time polynomial in n.*

Proof. To find a 2-bounded prefix from S to some vertex cover S' (or from T to some vertex cover T') we simply need to satisfy one of the conditions of Lemma 4, which can both be checked in time polynomial in n. Since $G[S_R \cup T_A]$ is a cactus graph, we know from Lemma 5 that one of them must be true. □

Theorem 3. Vertex Cover Reconfiguration *in cactus graphs can be solved in time polynomial in n.*

Proof. From Lemma 5, we know that for any cactus graph G and two vertex covers S and T of G, then either $S \overset{2,T}{\longleftrightarrow} S'$ or $T \overset{2,S}{\longleftrightarrow} T'$, where S' and T' are some vertex covers of G. In addition, Lemma 6 shows that such 2-bounded prefixes can be found in time polynomial in n. By combining these facts, we can now apply Lemma 3. That is, if $|S| \leq k-2$ and $|T| \leq k-2$, a reconfiguration sequence of length $|S_R| + |T_A|$ from S to T can be found in time polynomial in n.

When S (or T) has size k and is minimal, then we have a no-instance since neither removing nor adding a vertex results in a k-vertex cover, and hence S (or T) will be an isolated node in the reconfiguration graph, with no path to T (or S).

Table 1. Case Analysis. A ✓denotes a yes-instance and a ✗ denotes a no-instance.

		T non-minimal			T minimal		
		$k-2$	$k-1$	k	$k-2$	$k-1$	k
S non-minimal	$k-2$	✓	(1)	(5)	✓	(3)	✗
	$k-1$	(1)	(1)	(5)	(3)	(3)	✗
	k	(5)	(5)	(5)	(4)	(4)	✗
S minimal	$k-2$	✓	(3)	(4)	✓	(2)	✗
	$k-1$	(3)	(3)	(4)	(2)	(2)	✗
	k	✗	✗	✗	✗	✗	✗

The remaining cases to consider are listed in Table 1. Some of these cases are symmetric since the roles of S and T can be interchanged.

Case (1): $|S| = k-1$, $|T| = k-1$, S is non-minimal, and T is non-minimal. When both S and T are of size $k-1$ and are non-minimal, then each must contain at least one removable vertex. Hence, by removing such vertices, we can transform S and T into vertex covers S' and T', respectively, of size $k-2$. Finding a reconfiguration sequence of shortest length from S' to T' can be accomplished in time polynomial in n by Lemma 3. To guarantee a shortest path from S to T, we have to carefully select S' and T'. We let u and v be removable vertices in S and T respectively, and we set $S' = S \setminus \{u\}$ and $T' = T \setminus \{v\}$.

1. If $u \in S_R$ and $v \in T_A$, then the length of a shortest reconfiguration sequence from S' to T' will be $|S'_D| + |T'_A| = |S_R| + |T_A| - 2$. Therefore, accounting for the two additional removals, the length of a shortest path from S to T will be equal to $|S_R| + |T_A|$.
2. If $u \in S_R$ and $v \in C_{ST}$, then the length of a shortest reconfiguration sequence from S' to T' will be $|S'_D| + |T'_A| = |S_R| + |T_A| - 1$. Since v is in C_{ST}, it must be removed and added back. Therefore, the length of a shortest path from S to T will be equal to $|S_R| + |T_A| + 2$. The same is true when $u \in C_{ST}$ and $v \in T_A$ or when $u = v$ and $u \in C_{ST}$.
3. Otherwise, when $u \in C_{ST}$, $v \in C_{ST}$, and $u \neq v$, the length of a shortest path from S to T will be $|S_R| + |T_A| + 4$ since we have to touch two vertices in C_{ST} (i.e. two extra additions and two extra removals).

As there are at most $(k-1)^2$ pairs of removable vertices in S and T to check for Case (1), we can exhaustively try all pairs and choose one that minimizes the length of a reconfiguration sequence. Similar arguments can be applied when one of S or T is of size $k-2$, in which case we only need to check at most $k-1$ removable vertices.

Case (2): $|S| = k-1$, $|T| = k-1$, S is minimal, and T is minimal. If both S and T are minimal, of size $k-1$, and the minimum degree in $G[S_R \cup T_A]$ is two, then we have a no-instance; any removal will require at least two additions, therefore violating the maximum allowed capacity. If $G[S_R \cup T_A]$ has a vertex $u \in S_R$ such that $N_{G[S_R \cup T_A]}(u) = \{v\}$, we let S' be the vertex cover obtained by adding v and removing u as well as all vertices in $N_{G[S_R \cup T_A]}(v)$ that become removable after the addition of v. Since T is also minimal, we can apply the same transformation starting from T to get some vertex cover T'. These transformations correspond to 1-bounded prefixes. If both S' and T' are still minimal and of size $k-1$, we can exhaustively repeat this process until we either find a reconfiguration from S to T of length $|S_R| + |T_A|$, reach a state similar to Case (1),

or determine a no-instance (i.e. minimum degree two in $G[S'\Delta T']$). To see why Case (2) can be handled in time polynomial in n, we note that the described process is repeated at most $|S_R|+|T_A|$ times. At every iteration, we simply inspect the neighborhood of every vertex in the graph induced by the symmetric difference of the corresponding vertex covers. Finally, when one of S or T is of size $k-1$ and the other of size $k-2$, we only need to apply the transformations described above starting from the vertex cover of size $k-1$.

Case (3): $|S| = |T| = k-1$, S (T) is minimal, and T (S) is non-minimal. Case (3) can be handled by combining the arguments from Cases (1) and (2); we apply 1-bounded prefixes to the minimal vertex cover (Case (2)) and select a removable vertex from the non-minimal vertex cover which minimizes the total number of reconfiguration steps (Case (1)). Since $|S| = |T| = k-1$, we need to remove at most one vertex from S and one from T to obtain vertex covers of size $k-2$ to which we can apply Lemma 3 to obtain a reconfiguration sequence of shortest possible length. Hence, the length of a reconfiguration sequence from S to T will be at most $|S_R| + |T_A| + 4$, which occurs whenever we have to touch two vertices in C_{ST}. Whenever one of S or T is of size $k-2$, we only need to apply the described arguments to the vertex cover of size $k-1$.

Case (4): $|S| = k$, $|T| = k-1$, S is non-minimal, and T is minimal. Case (4) can be broken down into at most k instances which can be solved as in Cases (2) and (3). In each instance, we let $S' = S \setminus \{v\}$ where v is a removable vertex in S. If S' is minimal, we apply Case (2). Otherwise, we apply Case (3). Since $|S| = k$ and $|T| = k-1$, we need to remove at most 2 vertices from S and 1 from T to obtain vertex covers of size $k-2$ to which we can apply Lemma 3 to obtain a reconfiguration sequence of shortest possible length. Hence, the length of a reconfiguration sequence from S to T will be at most $|S_R| + |T_A| + 6$, which occurs whenever we have to touch three vertices in C_{ST}. Similarly to previous cases, whenever one of S or T is of size $k-2$, we only need to apply the described arguments to the vertex cover of size k.

Case (5): $|S| = k$, $|T| = k$, S is non-minimal, and T is non-minimal. Similarly, Case (5) can be broken down into at most k instances which can be solved as in Case (4). In each instance, we let $T' = T \setminus \{v\}$ where v is a removable vertex in T. The length of a reconfiguration sequence will be at most $|S_R| + |T_A| + 8$, which occurs whenever we have to touch four vertices in C_{ST}. When one of S or T is of size $k-2$ or $k-1$ and the other is of size k, Case (5) can be broken down into at most k instances which can be solved as in Case (3).

It is not hard to see that in all cases, Vertex Cover Reconfiguration in cactus graphs can be solved in time polynomial in n. This completes the proof. □

5 FPT Algorithm

In this section, we focus on Vertex Cover Reconfiguration on graphs of bounded degree. We start by showing that Vertex Cover Reconfiguration is **NP**-hard on graphs of degree at most d, for any $d \geq 4$, by proving **NP**-hardness on 4-regular graphs. The proof is based on the observation that the reconfiguration version of the problem is at least as hard as the compression version, defined as:

Vertex Cover Compression
Input: A graph $G = (V, E)$ and a vertex cover C of G such that $|C| = k \geq 1$
Parameter: k
Question: Does G have a vertex cover C' of size $k - 1$?

Next, we give an **FPT** algorithm for Vertex Cover Reconfiguration on graphs of bounded degree. Both the **NP**-hardness result and the **FPT** algorithm rely on the representation of reconfiguration sequences as nice edit sequences.

5.1 Compression Via Reconfiguration

Theorem 4. Vertex Cover Reconfiguration *is at least as hard as* Vertex Cover Compression.

Proof. We demonstrate by a reduction from the latter to the former. For (G, C, k) an instance of Vertex Cover Compression, we let $V(G) = \{v_1, \ldots, v_n\}$ and form $G' = (V_G \cup V_A \cup V_B, E_G \cup E')$, where G' consists of the disjoint union of a copy of G and a biclique $K_{k,k}$. Formally, we have:

$V_G = \{g_1, \ldots, g_n\}$
$V_A = \{a_1, \ldots, a_k\}$
$V_B = \{b_1, \ldots, b_k\}$
$E_G = \{\{g_i, g_j\} \mid g_i \in V_G, g_j \in V_G, \{v_i, v_j\} \in E(G)\}$
$E' = \{\{a_i, b_j\} \mid a_i \in V_A, b_j \in V_B, 1 \leq i \leq k, 1 \leq j \leq k\}$.

We let $(G', S, T, 3k - 1, 6k - 2)$ be an instance of Vertex Cover Reconfiguration, where $S = V_A \cup \{g_i \mid v_i \in C\}$ and $T = V_B \cup \{g_i \mid v_i \in C\}$. Clearly $|S| = |T| = 2k$ and both S and T are vertex covers of G'. We claim that G has a vertex cover of size $k - 1$ if and only if there is a reconfiguration sequence of length $6k - 2$ or less from S to T.

Before we can remove any vertex from V_A, we need to add all k vertices from V_B. But $2k + k = 3k > 3k - 1$, which violates the maximum allowed capacity. Therefore, if there is a reconfiguration sequence from S to T, then one of the vertex covers in the sequence must contain at most $2k - 1$ vertices. Of those $2k - 1$ vertices, k vertices correspond to the vertices in V_A and cover only the edges in E'. Thus, the remaining $k - 1$ vertices must be in V_G and should cover all the edges in E_G. By our construction of G', these $k - 1$ vertices correspond to a vertex cover of G.

Similarly, if G has a vertex cover C' such that $|C'| = k - 1$, then the following reconfiguration sequence transforms S to T: add all vertices of C', remove all vertices of C, add all vertices of V_B, remove all vertices from V_A, and finally add back all vertices of C and remove those of C'. The length of this sequence is equal to $6k - 2$ whenever $C \cap C' = \emptyset$ and is shorter otherwise. □

5.2 NP-Hardness on 4-Regular Graphs

We are now ready to show that Vertex Cover Reconfiguration remains **NP**-hard even if the input graph is restricted to be 4-regular. We use the same ideas as we did in the previous section. Since Vertex Cover remains **NP**-hard on 4-regular graphs [11] and any algorithm which solves the Vertex Cover Compression problem can be used to solve the Vertex Cover problem, we get the desired result.

The main difference here is that we need to construct a gadget, W_k, that is also 4-regular. We describe W_k in terms of several component subgraphs, each playing a role in forcing the reconfiguration of vertex covers.

A *k-necklace*, $k \geq 4$, is a graph obtained by replacing every edge in a cycle on k vertices by two vertices and four edges. For convenience, we refer to every vertex on the original cycle as a *bead* and every new vertex in the resulting graph as a *sequin*. The resulting graph has k beads each of degree four and $2k$ sequins each of degree two. Every two sequins that share the same neighborhood in a k-necklace are called a *sequin pair*. We say two beads are *adjacent* whenever they share exactly two common neighbors. Similarly, we say two sequin pairs are *adjacent* whenever they share exactly one common neighbor. Every two adjacent beads (sequin pairs) are *linked* by a sequin pair (bead).

The graph W_k consists of $2k$ copies of a k-necklace. We let $U = \{U_1, \ldots U_k\}$ and $L = \{L_1, \ldots L_k\}$ denote the first and second k copies respectively; for convenience, we use the terms "upper" and "lower" to mean "in U" and "in L", respectively. We let $b^u_{i,j}$ and $b^l_{i,j}$ denote the jth beads of necklace U_i and L_i respectively, where $1 \leq i \leq k$ and $1 \leq j \leq k$. Beads on each necklace in W_k are numbered consecutively in clockwise order from 1 to k. For every two adjacent beads $b^x_{i,j}$ and $b^x_{i,j+1}$, where $x \in \{u, l\}$, we let $\{p^x_{i,j}\}$ denote the sequin pair which links both beads.

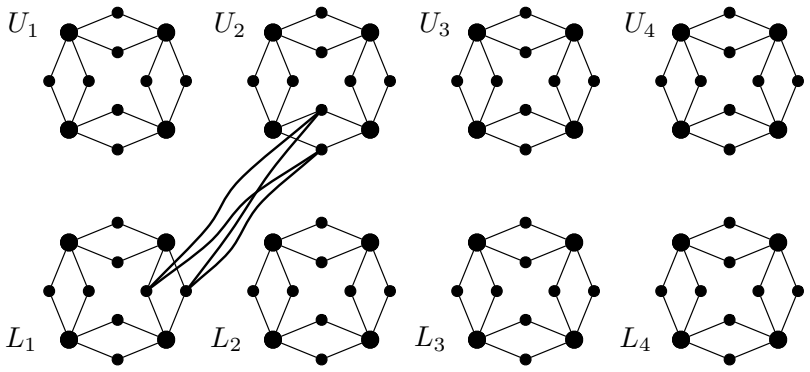

Fig. 1. The graph W_4 (the edges of one out of the k^2 joining bicliques is shown)

For each sequin pair $p^l_{i,j}$, we add four edges to form a $K_{2,2}$ (a *joining biclique*) with the pair $p^u_{j,i}$, for all $1 \leq i, j \leq k$ (Figure 1); we say that sequin pairs $p^l_{i,j}$ and $p^u_{j,i}$ are *joined*. All k^2 joining bicliques in W_k are vertex disjoint. The total number of vertices in W_k is $6k^2$. Every vertex has degree exactly four; every bead is connected to four sequins from the same necklace and every sequin is connected to two beads from the same necklace and two other sequins from a different necklace. We let S be the set containing all upper beads and lower sequins, whereas T contains all lower beads and upper sequins. Formally, $S = \{b^u_{i,j} \mid 1 \leq i, j \leq k\} \cup \{v \in p^l_{i,j} \mid 1 \leq i, j \leq k\}$ and $T = \{b^l_{i,j} \mid 1 \leq i, j \leq k\} \cup \{v \in p^u_{i,j} \mid 1 \leq i, j \leq k\}$. Each set contains $3k^2$ vertices, that is, half the vertices in W_k.

Observation 11 *S and T are minimum vertex covers of W_k.*

Proof. We need at least $2k^2$ vertices to cover the edges in the k^2 vertex disjoint joining bicliques contained in W_k. Moreover, any minimal vertex cover C of W_k which includes a vertex v from a sequin pair $p^x_{i,j} = \{v, w\}$, where $x \in \{u, l\}$, must also include w. Otherwise, the two beads linking

$p^x_{i,j}$ to its adjacent sequin pairs must be in C to cover the edges incident on w, making v removable. Hence, any minimal vertex cover C of W_k must include either one or both sequin pairs in a joining biclique. We let x denote the number of joining bicliques from which two sequin pairs are included in C. Similarly, we let y denote the number of joining bicliques from which only one sequin pair is included in C. Hence, $x + y = k^2$ and $|C| \geq 4x + 2y$. When $y = 0$, $|C| \geq 4k^2$ and C cannot be a minimum vertex cover, as S and T are both vertex covers of W_k of size $3k^2$. When $y \geq 1$, we are left with at least y uncovered edges incident to the sequin pairs not in C. Those edges must be covered using at least y beads and hence $|C| \geq 4x + 3y$. If we assume $4x + 3y < 3k^2$, we get a contradiction since $4x + 4y = 4k^2 < 3k^2 + y$ and $k^2 < y$. Therefore, S and T must be minimum vertex covers of W_k. □

To prove the next two lemmas, we consider the representation of reconfiguration sequences as nice edit sequences. We know from Lemma 1 that any reconfiguration sequence of length $6k^2$ from S to T can be converted into a nice edit sequence α of the same length. Since S is a minimal vertex cover of W_k, α cannot start with a vertex removal and hence the starting piece of α must be empty. Since $V(S, \alpha[1, |\alpha| - 1])$ is a vertex cover of W_k, $|S| = |T|$, and S and T are minimum vertex covers of W_k, α cannot end with a vertex addition and hence the ending piece of α must also be empty. Moreover, $|\alpha| = 6k^2$ implies that α must touch every vertex in W_k exactly once. Since W_k is a 4-regular graph, each d-add-remove segment in the central piece of α consists of at most four additions followed by at most four removals. Hence, we know $\alpha = \alpha_1\alpha_2 \ldots \alpha_j$, where each α_i is a 4-add-remove segment in α.

Lemma 7. *Any nice edit sequence α of length $6k^2$ in W_k from S to T either adds or removes both vertices u and v in a sequin pair in the same 4-add-remove segment β.*

Proof. Both vertices in a sequin pair share the same neighborhood. Hence, when u is removed, all of its neighbors must have been added, making v also removable. Moreover, since every vertex is touched exactly once in α, none of the neighbors of u and v will be touched in α after the removal of u. Therefore, if v is not removed in the same segment β as u, the early removal invariant of Definition 3 will be violated since v is separated from any vertex that gets touched after β.

For the case of additions, if only u is added in β then none of its neighbors can be removed. It follows that u is not connected to any vertex in $V(\beta)$, hence violating the connectivity invariant of Definition 3. □

Lemma 8. *There exists a function of k, $f(k)$, such that $(W_k, S, T, 3k^2 + f(k), \ell)$ is a yes-instance and $(W_k, S, T, 3k^2 + f(k) - 1, \ell)$ is a no-instance of* Vertex Cover Reconfiguration *for $\ell = 6k^2$. Moreover, $k - 2 \leq f(k) \leq k + 3$.*

Proof. To show that such an $f(k)$ exists, we first prove the $k - 2$ lower bound by showing that any nice edit sequence $\alpha = \alpha_1\alpha_2 \ldots \alpha_j$ of length $6k^2$ from S to T must have some prefix with exactly $5k$ vertex removals and at least $6k - 2$ additions. We let position x, $1 \leq x \leq |\alpha|$ be the smallest position such that $\alpha[1, x]$ contains exactly $5k$ vertex removals. Those $5k$ vertices correspond to a set $S' \subset S$, as α touches every vertex exactly once. The claim is that $\alpha[1, x]$ must contain at least $6k-2$ vertex additions. We let $T' \subset T$ denote the set of added vertices in $\alpha[1, x]$. Since $N_{W_k}(S') \subseteq T'$, we complete the proof of the lower bound by showing that $|T'| \geq |N_{W_k}(S')| \geq \frac{6}{5}|S'| - 2 \geq 6k - 2$. To do so, we show that for any $S' \subset S$ of size $5k$, $N_{W_k}(S') \subseteq T'$ contains at least $\frac{6}{5}|S'| - 2 = 6k - 2$ vertices.

In what follows, we restrict our attention to the bipartite graph $Z = W_k[S' \cup T']$ and we let S' and T' denote the two partitions of Z. We subdivide S' into two sets: S'_b contains upper beads and S'_s contains lower sequins. Since every vertex in S'_b has four neighbors in T' and adjacent beads share exactly two neighbors, we have $|N_Z(S'_b)| \geq 2|S'_b|$ and equality occurs whenever S'_b contains $2k$ beads from the same two upper necklaces. Whenever S'_b contains less than $2k$ beads and $Z[S'_b \cup N_Z(S'_b)]$ consists of $t_b \geq 1$ connected components, at least one bead from each component (except possibly the first) will be adjacent to at most one other bead in the same component. Therefore, $|N_Z(S'_b)| \geq 2|S'_b| + 2(t_b - 1)$.

Lemma 7 implies that T' will always contain both vertices of any sequin pair. Since we are only considering vertices in $V(\alpha[1, x])$, some sequins in S'_s might be missing the other sequin in the corresponding pair. However, all the neighbors of the sequin pair have to be in T' so we assume without loss of generality that vertices in S'_s can be grouped into sequin pairs. Every sequin pair in S'_s has four neighbors in T'. Adjacent sequin pairs share exactly one neighbor. Hence, $|N_Z(S'_s)| \geq \frac{3}{2}|S'_s|$ and equality occurs whenever S'_s contains k sequin pairs of a single lower necklace. Whenever S'_s contains less than k sequin pairs and $Z[S'_s \cup N_Z(S'_s)]$ consists of $t_s \geq 1$ connected components, at least one sequin pair from each component will be adjacent to at most one other sequin pair in the same component. Therefore, $|N_Z(S'_s)| \geq \frac{3}{2}|S'_s| + t_s$.

Combining the previous observations, we know that when either S'_b or S'_s is empty, $|N_Z(S')| \geq \frac{6}{5}|S'|$, as needed. When both are not empty, we let $I = N_Z(S'_b) \cap N_Z(S'_s)$. Hence, $|N_Z(S'_b)| + |N_Z(S'_s)| - |I| \geq 2|S'_b| + 2(t_b - 1) + \frac{3}{2}|S'_s| + t_s - |I|$ and we rewrite it as:

$$|N_Z(S')| + 2 \geq \frac{100}{50}|S'_b| + \frac{75}{50}|S'_s| + 2(t_b - 1) + t_s - (|I| - 2) \tag{8}$$

We now bound the size of I. Note that I can only contain upper sequin pairs joined with sequin pairs in S'_s. As every sequin pair in S'_s has either zero or two neighbors in I, $|S'_s| \geq |I|$. Moreover, for every two sequin pairs in S'_s having two neighbors in I, there must exist at least one vertex in S'_b, which implies $|S'_b| \geq \frac{|I|}{4}$. Finally, whenever a sequin pair $p \in S'_s$ has two neighbors in I, then $t_b, t_s \geq 1$ as at least one bead neighboring the sequin pair joined with p must be in S'_b. Every other sequin pair $p' \in S'_s$, $p' \neq p$, with two neighbors in I will force at least one additional connected component in either $Z[S'_b \cup N_Z(S'_b)]$ or $Z[S'_s \cup N_Z(S'_s)]$ since W_k contains a single joining biclique between any two necklaces. Therefore, the total number of connected components is $t_b + t_s \geq \frac{|I|}{2}$. Putting it all together, we get:

$$\begin{aligned} \frac{40}{50}|S'_b| + \frac{15}{50}|S'_s| + 2(t_b - 1) + t_s &\geq \frac{2}{10}|I| + \frac{3}{10}|I| + \frac{5}{10}|I| + t_b - 2 \\ &\geq |I| - 2 \end{aligned} \tag{9}$$

Combining Equations 8 and 9, we get:

$$\begin{aligned} |N_Z(S')| + 2 &\geq \frac{6}{5}|S'| + \frac{40}{50}|S'_b| + \frac{15}{50}|S'_s| + 2(t_b - 1) + t_s - (|I| - 2) \\ &\geq \frac{6}{5}|S'| \end{aligned} \tag{10}$$

Therefore, $V(S, \alpha[1, x])$ is a vertex cover of W_k of size at least $3k^2 + k - 2$, as needed.

To show the $f(k) \leq k+3$ upper bound, we show that $(W_k, S, T, 3k^2+k+3, 6k^2)$ is a yes-instance by providing an actual reconfiguration sequence (that is not nice):

(1) Add all k beads in L_1. Since S is a vertex cover of W_k, we know that the additional k beads will result in a vertex cover of size $3k^2 + k$.
(2) Add both vertices in $p^u_{1,1}$ and remove both vertices in $p^l_{1,1}$. The removal of both vertices in $p^l_{1,1}$ is possible since we added all their neighbors in L_1 (step (1)) and U_1. The size of a vertex cover reaches $3k^2 + k + 2$ after the additions and then drops back to $3k^2 + k$.
(3) Repeat step (2) for all sequin pairs $p^u_{i,1}$ and $p^l_{1,i}$ for $2 \leq i \leq k$. The size of a vertex cover is again $3k^2 + k$ once step (3) is completed. Step (2) is repeated a total of k times. After every repetition, we have a vertex cover of W_k since all beads in L_1 were added in step (1) and the remaining neighbors of each sequin pair in U_i are added prior to the removals.
(4) Add both vertices in $p^u_{1,2}$ and remove vertex $b^u_{1,2}$.
(5) Add $b^l_{2,1}$ and $b^l_{2,2}$. At this point, the size of a vertex cover is $3k^2 + k + 3$.
(6) Remove both vertices in $p^l_{2,1}$.
(7) Repeat steps (4), (5), and (6) until all beads in L_2 have been added and the sequin pairs removed. When we reach the last sequin pair in L_2, $b^l_{2,1}$ was already added and hence we gain a surplus of one which brings the vertex cover size back to $3k^2 + k$.
(8) Repeat steps (4) to (7) for every remaining necklace in L.

Since every vertex in W_k is touched exactly once, we know that $\ell = 6k^2$. In the course of the described reconfiguration sequence, the maximum size of any vertex cover is $3k^2 + k + 3$. Hence, $f(k) \leq k + 3$. This completes the proof. □

It would be interesting to close the gap on $f(k)$, but the existence of such a value is enough to prove the main theorem of this section.

Theorem 5. Vertex Cover Reconfiguration *is* **NP***-hard on* 4*-regular graphs.*

Proof. We demonstrate by a reduction from Vertex Cover Compression to Vertex Cover Reconfiguration where the input graph is restricted to be 4-regular in both cases. For (G, C, k) an instance of Vertex Cover Compression, we form $G' = (V(G) \cup V(W_k), E(G) \cup E(W_k))$. We let $(G', S, T, 3k^2+k+f(k)-1, 6k^2+4k-2)$ be an instance of Vertex Cover Reconfiguration, where $S = \{e^u_{i,j} \mid 1 \leq i, j \leq k\} \cup \{p^l_{i,j} \mid 1 \leq i, j \leq k\} \cup C$ and $T = \{e^l_{i,j} \mid 1 \leq i, j \leq k\} \cup \{p^u_{i,j} \mid 1 \leq i, j \leq k\} \cup C$ and $f(k)$ is the value whose existence was shown in Lemma 8.

Clearly $|S| = |T| = 3k^2 + k$ and both S and T are vertex covers of G'. We claim that G has a vertex cover of size $k-1$ if and only if there is a reconfiguration sequence of length $6k^2 + 4k - 2$ or less from S to T.

We know from Lemma 8 that the reconfiguration of W_k requires at least $f(k)$ available capacity. But $3k^2 + k + f(k) > 3k^2 + k + f(k) - 1$, which violates the maximum allowed capacity. Therefore, if there is a reconfiguration sequence from S to T, then one of the vertex covers in the sequence must contain at most $3k^2+k-1$ vertices. By Observation 11, we know that $3k^2$ of those $3k^2+k-1$ vertices are needed to cover the edges in $E(W_k)$. Thus, the remaining $k-1$ vertices must be in $V(G)$ and should cover all the edges in $E(G)$. By our construction of G', these $k-1$ vertices correspond to a vertex cover of G.

Similarly, if G has a vertex cover C' such that $|C'| = k-1$, then the following reconfiguration sequence transforms S to T: add all vertices of C', remove all vertices of C, apply the reconfiguration

sequence whose existence was shown in Lemma 8 to $G'[V(W_k)]$, and finally add back all vertices of C and remove those of C'. The length of this sequence is equal to $6k^2+4k-2$ whenever $C \cap C' = \emptyset$ and is shorter otherwise. □

5.3 FPT Algorithm for Graphs of Bounded Degree

In this section, we prove the following theorem:

Theorem 6. Vertex Cover Reconfiguration *parameterized by ℓ is fixed-parameter tractable for graphs of degree at most d.*

Before we discuss the technical details, we give a few additional definitions, observations, and lemmas. We also introduce a new problem, which we will use to develop our **FPT** algorithm. We extensively use the different kinds of edit sequences defined in Section 2.1. We sometimes use the modified big-Oh notation $\mathcal{O}^*$ that suppresses all polynomially bounded factors.

The following observations will be useful for determining an upper bound on the running time of our algorithm. Given a (labeled) graph G, we say two subgraphs of G are *vertex-differing* whenever they differ in at least one vertex. Recall that for $r \geq 0$, the r-neighborhood of a vertex $v \in V(G)$ is defined as $N^r_G[v] = \{u \mid dist_G(u,v) \leq r\}$, $B(v,r) = N^r_G[v]$ is a ball of radius r around v and for $A \subseteq V(G)$, $B(A,r) = \bigcup_{v \in A} N^r_G[v]$.

Observation 12 *For any graph G of degree at most d, $v \in V(G)$, and $A \subseteq V(G)$, $|B(v,r)| \leq d^{r+1}$ and $|B(A,r)| \leq |A|d^{r+1}$.*

Observation 13 *Given a graph G of degree at most d, every vertex $v \in V(G)$ can appear in at most $d^{s^2(r+1)}$ vertex-differing connected subgraphs of G having at most s vertices and diameter at most r.*

Proof. Since G has degree at most d and the diameter of every connected subgraph must be at most r, the number of vertices which can belong to the same connected subgraph as v is at most d^{r+1} (Observation 12). Of those d^{r+1} vertices, there are at most $\sum_{i=1}^{s} \binom{d^{r+1}}{i} < d^{s^2(r+1)}$ possible connected subgraphs on at most s vertices which include v and have diameter at most r. □

Observation 14 *Given a set S of vertices, $|S| \geq r$, the number of ways we can order at most r vertices from S into a sequence of size exactly r is less than $|S|^{2r}$.*

Proof. There are $\binom{|S|}{r}$ possible subsets of size r in S. For each subset, there are at most r^r ways we can order the vertices into a sequence of size exactly r. Since $|S| \geq r$, we get the desired bound. □

Observation 15 *Given a set S of vertices, the total number of possible full labeled edit sequences of size at most ℓ touching only vertices in S is at most $2^{\ell+1}|S|^{2\ell+2}$. The total number of possible partial labeled edit sequences of size at most ℓ touching only vertices in S is at most $3^{\ell+1}|S|^{2\ell+2}$.*

Proof. There are at most 2^r possible full unlabeled edit sequences of size exactly r. In addition, Observation 14 states that there are at most $|S|^{2r}$ ways we can order at most r vertices from S into a sequence of size exactly r. Therefore, combining both observations, we get the desired bound of at most $\sum_{i=1}^{\ell} 2^i|S|^{2i} < 2^{\ell+1}|S|^{2\ell+2}$ possible full labeled edit sequences of size at most ℓ. Using similar arguments, we get the $\sum_{i=1}^{\ell} 3^i|S|^{2i} < 3^{\ell+1}|S|^{2\ell+2}$ bound on the total number of possible partial labeled edit sequences of size at most ℓ. □

We are now ready to introduce the Vertex Cover Walk (VCW) problem, which is formally defined as follows:

Vertex Cover Walk
Input: A graph G, a vertex cover S of G, and a full unlabeled edit sequence α of size $\ell \geq 1$
Parameter: $|\alpha| = \ell$
Question: Can we apply α to G and S?

Using an algorithm for the VCW problem, one can easily solve the ℓ-Independent Set problem, i.e., finding an independent set of size ℓ in a graph $G = (V, E)$, by simply setting $S = V$ and $\alpha = r^\ell$, where r^ℓ is an unlabeled edit sequence consisting of ℓ removal markers. In fact, we can show that, in the parameterized setting, VCW is at least as hard as Vertex Cover Local Search (VCLS) [8], defined as:

Vertex Cover Local Search
Input: A graph G, a vertex cover S of G, and integer $\ell \geq 1$
Parameter: ℓ
Question: Does G have a vertex cover S' such that $|S'| < |S|$ and $|S' \Delta S| \leq \ell$?

Lemma 9. Vertex Cover Walk *is at least as hard as* Vertex Cover Local Search

Proof. By Observation 3, any algorithm which solves the VCW problem in $\mathcal{O}(f(\ell))$ time, for some computable function f, can solve the VCLS problem in $\mathcal{O}(f(\ell)2^{\ell+1})$ time by simply enumerating all possible full unlabeled edit sequences of size at most ℓ where the number of removals is at least one greater than the number of additions. □

VCLS is known to be $\mathbf{W}[\mathbf{1}]$-hard on graphs of bounded degeneracy and fixed-parameter tractable on graphs of bounded degree [8]. Finally, we also make use of the ℓ-Multicolored Independent Set (ℓ-MIS) problem:

ℓ-Multicolored Independent Set
Input: A graph G, a positive integer ℓ, and a vertex-coloring $c : V(G) \to \{c_1, \ldots, c_\ell\}$ for G
Parameter: ℓ
Question: Does G have an independent set of size ℓ including exactly one vertex of each color?

Note that edges in $E(G)$ need not have endpoints assigned different colors. The ℓ-MIS problem is $\mathbf{W}[\mathbf{1}]$-hard in general graphs as we can reduce the well-known $\mathbf{W}[\mathbf{1}]$-hard ℓ-Multicolored Clique problem to it by simply complementing all edges in the input graph [7]. For a vertex $v \in V(G)$, we denote by $c(v)$ the color assigned to v. We let $V_i(G)$ denote the set of vertices assigned colored c_i in G, i.e. $V_i(G) = \{v \in V(G) \mid c(v) = c_i\}$. We say vertex v belongs to *color class* c_i if $c(v) = c_i$.

Lemma 10. *The* ℓ-Multicolored Independent Set *problem is fixed-parameter tractable if for every vertex* $v \in V(G)$ *such that* $c(v) = c_i$, $|N_G(v) \cap V_j(G)| \leq d$, *for some fixed integer* d, $i \neq j$, *and* $1 \leq i, j \leq \ell$. *Moreover, there is an algorithm which solves the problem in* $\mathcal{O}^*((d\ell)^{2\ell})$ *time.*

Proof. Since every vertex $v \in V(G)$ assigned color c_i has at most d neighbors assigned color c_j for all $j \neq i$, v has at most $d\ell$ neighbors not in $V_i(G)$.

If $|V_i(G)| > d\ell$ for all i between one and ℓ, we can simply pick $d\ell + 1$ vertices from each $V_i(G)$ and delete all remaining vertices from G to obtain G'. By exhaustively enumerating all possible subsets of ℓ vertices assigned different colors in G', we are guaranteed to find an independent set of size ℓ. For v_1 a vertex in $V_1(G)$, since v_1 has at most d neighbors in $V_2(G)$, there must exist a vertex $v_2 \in V_2(G)$ such that $\{v_1, v_2\} \notin E(G')$. Similarly, v_1 and v_2 have at most d neighbors in $V_3(G)$, for a total of at most $2d$ neighbors. Since $|V_3(G)| = d\ell + 1 > 2d$, there must exist a vertex $v_3 \in V_3(G)$ such that $\{v_1, v_3\}, \{v_2, v_3\} \notin E(G')$. By repeating the same argument, we know that once we reach $V_\ell(G)$, all the previously selected vertices can have at most $(\ell - 1)d < d\ell + 1$ distinct neighbors in $V_\ell(G)$. Therefore, we know that there must exist an independent set of size ℓ in G' which includes exactly one vertex of each color. Moreover, we can exhaust all possibilities in $\mathcal{O}^*(\binom{d\ell^2+\ell}{\ell})$ time.

Otherwise, we know that the size of some $V_i(G)$ is at most $d\ell$. In this case, we can generate $d\ell$ new instances of the problem, where in each instance we delete all but one of the vertices in $V_i(G)$ to obtain G'. If G' has some $V_i(G)$ with at most $d\ell$ vertices, we repeat the process. Otherwise, we apply the exhaustive search described above. It is not hard to see that this algorithm runs in $\mathcal{O}^*((d\ell)^{2\ell})$ time. □

High Level Description

Our algorithm relies on a combination of brute-force enumeration and reductions to the problems VCW and ℓ-MIS. We make use of instances of Annotated VCR (AVCR) in which the vertex set of the input graph is partitioned into sets X, W, and R such that, X and R are separated, $S_R \cup T_A \subseteq X$, and no vertex in W is touched during reconfiguration.

Annotated VCR
Input: A graph $G = (V, E)$, positive integers k and ℓ,
two vertex covers S and T of G of size at most k,
and partition X, W, and R of $V(G)$.
Parameter: ℓ
Question: Is there a reconfiguration sequence from S to T
of length at most ℓ touching no vertex in W?

The algorithm implicit in the proof of Lemma 11 generates 2ℓ instances $\{\mathcal{I}_1, \ldots, \mathcal{I}_{2\ell}\}$ of AVCR such that the original instance is a yes-instance of VCR if and only if some $\mathcal{I}_x$ is a yes-instance of AVCR (Lemma 11). In each instance $\mathcal{I}$, solved by the Solve-AVCR algorithm (Algorithm 2), there is a subset W of C_{ST} separating a superset X of $S_R \cup T_A$ from a vertex set R such that no vertex in W is touched during reconfiguration. We use brute-force enumeration to generate all partial labeled edit sequences that touch only vertices in X; if any produces a tight sequence that transforms S to T, we have a yes-instance. Otherwise, we consider all such sequences which are valid and transform S to T but exceed the capacity constraint. By finding an appropriate labeled filling sequence γ' that touches vertices in R, we can free up capacity so that $merge(\beta, \gamma')$ is tight.

We can find such a γ' trivially if there is a sufficiently large independent set in the vertices of $S \cap R$ with no neighbours in O_{ST}, as our reconfiguration sequence will consist of removing the vertices in the independent set to free up capacity, applying β, and then adding back the vertices in the independent set. Otherwise, we reduce the problem of finding γ' to one instance of VCW for each suitable unlabeled edit sequence γ of size the number of blanks in β. The walk must correspond

to a walk from the node corresponding to $S \cap R$ in $R_{\mathrm{VC}}(G[R], 0, k)$ back to the node itself. Any labeled edit sequence γ' returned by the Solve-VCW algorithm (Algorithm 3) will be nice, and hence a yes-instance if $merge(\beta, \gamma')$ is tight and transforms S into T.

To try all possible ways of generating γ', we start by enumerating all d-well-formed full unlabeled edit sequences γ of the appropriate size, and for each try all possible choices for the starting piece; as any reconfiguration sequence can be converted to a nice one, it suffices to restrict our examination to nice sequences (and correspondingly, d-well-formed sequences without labels). Given γ and a starting piece, we create t instances of VCW, where instance $\mathcal{J}_y$ corresponds to the graph induced by the labeled central piece having y connected components. To solve instance $\mathcal{J}_y$, we consider all ways of assigning the d-add-remove segments to connected components. This allows us to create a sequence for each component, where γ_h touches only vertices in component h, and all vertices touched by γ'_h can be found in a ball of radius $|\gamma'_h|$.

We can reduce each subproblem to an instance of y-MIS, where a color c_h corresponds to component h. In our auxiliary graph G_A (Algorithm 4), we create vertices for each labeled full edit sequence λ, where $G[V(\lambda)]$ is connected and λ can be derived by adding labels to γ_h. There is an edge between two vertices in G_A if they have different colors and sets of vertices associated with their edit sequences are not separated. Thus, a solution to y-MIS indicates that there are y full labeled edit sequences with separated vertex sets, as required to complete the central piece of γ. As the ending piece of γ' will be determined by the starting and central pieces, this completes the algorithm.

Technical Details

In this last section, we discuss the technical details of our algorithms and give a bound on the overall running times. We make use of *Ramsey Numbers*, i.e. for any positive integers p and q, there exists a minimum number $\mathcal{R}(p, q)$ (Ramsey Number), such that any graph on at least $\mathcal{R}(p, q)$ vertices contains either a clique of size p or an independent set of size q. Moreover, $\mathcal{R}(p, q) \leq \binom{p+q-2}{q-1}$ [14].

Lemma 11. *For any instance of* Vertex Cover Reconfiguration*, there exists a set of* 2ℓ *instances* $\{\mathcal{I}_1, \ldots, \mathcal{I}_{2\ell}\}$ *of* Annotated VCR *such that the original instance is a yes-instance for* VCR *if and only if at least one* $\mathcal{I}_x$ *is a yes-instance for* AVCR*,* $1 \leq x \leq 2\ell$ *and for each* X_x*,* $S_R \cup T_A \subseteq X_x$*.*

Proof. Recall that for an instance (G, S, T, k, ℓ) of VCR, we partition $V(G)$ into the sets $C_{ST} = S \cap T$, $S_R = S \setminus C_{ST}$, $T_A = T \setminus C_{ST}$, and the independent set $O_{ST} = V(G) \setminus (S \cup T) = V(G) \setminus (C_{ST} \cup S_R \cup T_A)$. We further subdivide some of these sets (Figure 2).

We let C_i be the set of vertices in C_{ST} that are at distance i from $S_R \cup T_A$. That is, we set $C_i = \{B(S_R \cup T_A, i) \setminus B(S_R \cup T_A, i-1)\} \cap C_{ST}$, for $i \geq 1$. By Observation 1, there can be no edges between vertices in $S_R \cup T_A$ and vertices in O_{ST}. Therefore, we let $O_i = \{B(S_R \cup T_A, i) \setminus B(S_R \cup T_A, i-1)\} \cap O_{ST}$ be the set of vertices in O_{ST} that are at distance i from $S_R \cup T_A$, for $i \geq 2$. We let C_∞ and O_∞ be the sets of vertices in C_{ST} and O_{ST}, respectively, which are not in the r-neighborhood of $S_R \cup T_A$ for any value of r. Note that $C_\infty \cup O_\infty$ is not necessarily connected.

By our definition of C_i, every vertex $v \in C_i$, $i \geq 1$, must be at distance i from some vertex in $S_R \cup T_A$. Similarly, every vertex $v \in O_i$, $i \geq 2$, must be at distance i from some vertex in $S_R \cup T_A$. Therefore, since O_i is an independent set, for any vertex $v \in O_i$, $i \geq 2$, all the neighbors of v must be in $C_{i-1} \cup C_i \cup C_{i+1}$. Since there are no edges between O_{ST} and $S_R \cup T_A$, the set C_1 separates $S_R \cup T_A$ from the rest of the graph. Moreover, for any $i \geq 2$, the vertices in $C_i \cup C_{i+1}$ separate

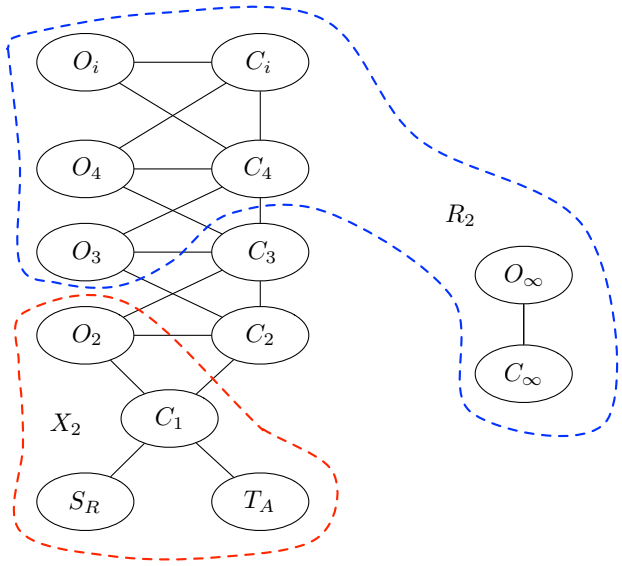

Fig. 2. The subdivision of the sets C_{ST} and O_{ST} along with the sets X_2 and R_2.

$X_i = S_R \cup T_A \cup \bigcup_{1 \leq j < i} C_j \cup \bigcup_{2 \leq j \leq i} O_j$ from the rest of the graph $R_i = V(G) \setminus \{X_i \cup C_i \cup C_{i+1}\}$ (Figure 2). We say $W_i = V(G) \setminus \{X_i \cup R_i\}$ is a *wall-set*. The separation of X_i from R_i plays a crucial role in our algorithm. Since X_i is a subset of $B(S_R \cup T_A, i)$, it follows from Observation 12 that:

Observation 16 $|X_i| \leq |S_R \cup T_A| d^{i+1} \leq \ell d^{i+1}$.

In the first instance $\mathcal{I}_1$, the inputs to the Solve-AVCR algorithm consist of the sets $X = S_R \cup T_A$, $W = C_1$, and $R = V(G) \setminus \{X_1 \cup C_1\}$. For the $2\ell - 1$ remaining instances, $X = X_x$, $W = C_x \cup C_{x+1}$, and $R = R_x$, for $2 \leq x \leq 2\ell$. In all instances, the vertices in X are separated from the vertices in R. This separation allows us to "divide" the problem into two "subproblems", where the first can be solved by brute-force enumeration (Algorithm 2, lines 3–11) and the second can be solved via the Solve-VCW algorithm (Algorithm 2, lines 12–18). In each instance $\mathcal{I}_x$, $1 \leq x \leq 2\ell$, we force the vertices in the wall-set W to remain in every vertex cover throughout any reconfiguration sequence, i.e. we do not allow any of those vertices to be touched. Recall that all vertices in the wall-set must be in C_{ST}. Assume that the Solve-AVCR algorithm (Algorithm 2) can solve each of those 2ℓ instances in **FPT** time without touching the vertices in the wall-set. If any of those instances is a yes-instance we are done, since we have found a reconfiguration sequence of length at most ℓ from S to T. Otherwise, we know from the first no-instance that if a reconfiguration sequence of length at most ℓ from S to T exists, then it must touch some vertex in $W = C_1$. Generally, we know from the xth no-instance, for $x > 1$, that if a labeled reconfiguration sequence of length at most ℓ from S to T exists, then it must touch some vertex in $W = C_x \cup C_{x+1}$. When all 2ℓ instances are no-instances, we know that any reconfiguration sequence from S to T must touch a vertex from each of the following sets: $C_1, C_2 \cup C_3, C_3 \cup C_4, \ldots, C_{2\ell} \cup C_{2\ell+1}$. However, $C_1 \cap \{C_2 \cup C_3\} \cap \{C_4 \cup C_5\} \cap \ldots \cap \{C_{2\ell} \cup C_{2\ell+1}\} = \emptyset$ and therefore at least ℓ vertices from C_{ST} must be touched, which implies that our original VCR instance is a no-instance. □

Algorithm 2 SOLVE-AVCR

Input: A graph G of degree at most d, a positive integer k, two vertex covers S and T of G of size at most k, a positive integer ℓ, and partition X, W, and R of $V(G)$.
Output: A reconfiguration sequence of length at most ℓ from S to T if one exists and $\varnothing$ otherwise.

1: Enumerate every partial labeled edit sequence β of size at most ℓ that only touches vertices in X;
2: **for each** β
3: **if** (β is tight and transforms S into T)
4: **return** $clean(\beta)$;
5: **if** (β is valid and transforms S into T and $cap(clean(\beta)) - k \leq \ell$)
6: $c = cap(clean(\beta)) - k$;
7: $S' = \{S \cap R\} \setminus \{v \mid v \in S \cap R \ \wedge |N_G(v) \cap O_{ST}| = 0\}$;
8: **if** ($|S'| > \mathcal{R}(d+2, c)$)
9: Find an independent set of size c in S';
10: Let L be the set of vertex labels (in any order) in the independent set;
11: **return** $concat(concat(label(r^c, L), clean(\beta)), label(a^c, L))$;
12: Generate every d-well-formed unlabeled γ of size $|\beta| - |clean(|\beta|)|$;
13: **for each** γ
14: $\gamma' =$ SOLVE-VCW$(G[R], S \cap R, \gamma)$;
15: **if** (γ' touches some vertex an odd number of times)
16: **continue**;
17: **if** ($merge(\beta, \gamma')$ is tight and transforms S into T)
18: **return** $merge(\beta, \gamma')$;
19: **return** $\varnothing$;

Lemma 12. *If* VERTEX COVER WALK *is solvable in $\mathcal{O}^*(f(d,\ell))$ time, for some computable function f, on a graph G of degree at most d, then* VERTEX COVER RECONFIGURATION *is solvable in $\mathcal{O}^*(2\ell 3^{\ell+1}(\ell d^{2\ell+1})^{2\ell+2}(d+\ell)^{\ell} 2^{\ell} f(d,\ell))$ time on G.*

Proof. Assuming an algorithm which solves the VCW in $\mathcal{O}^*(f(d,\ell))$ time, the worst case running time of the SOLVE-AVCR algorithm (Algorithm 2) is in $\mathcal{O}^*(3^{\ell+1}(\ell d^{2\ell+1})^{2\ell+2}\ (d+\ell)^{\ell}\ 2^{\ell}\ f(d,\ell))$; the size of X is bounded above by $\ell d^{2\ell+1}$ in the worst case (Observation 16). Hence, we have a total of at most $3^{\ell+1}(\ell d^{2\ell+1})^{2\ell+2}$ partial labeled edit sequences to enumerate (Observation 15). Since $G[S']$ has degree at most d, it cannot contain a clique of size $d+2$. Hence, if $|S'| > \mathcal{R}(d+2, c)$ (Algorithm 2, line 8), then $G[S']$ must contain an independent set of size $c < \ell$, which we can find in $\mathcal{O}^*(\binom{R(d+2,\ell)}{\ell})$ or $\mathcal{O}^*((d+\ell)^{\ell})$ time [20]. When $|S'| \leq \mathcal{R}(d+2, c)$, the full unlabeled edit sequences we consider must be of size at most $\ell - 1$. By Observation 3, there are at most 2^{ℓ} such sequences to check for each β (Algorithm 2, lines 12–18). Merging, cleaning, and concatenating sequences as well as checking whether they are valid or tight can be accomplished in time polynomial in the size of the sequences and n. Since there are at most 2ℓ instances of AVCR to solve, we get the desired bound. □

Lemma 13. *The* SOLVE-VCW *algorithm (Algorithm 3) returns a full labeled edit sequence $label(\gamma, L)$ that is nice starting at $S \cap R$, if one exists.*

Proof. The inputs to the SOLVE-VCW algorithm will consist of the graph $G[R]$, the vertex cover $S \cap R$ of G, and a d-well-formed full unlabeled edit sequence γ. We let γ_s, γ_c, and γ_e denote the starting, central, and ending pieces of γ respectively. Moreover, we let $\gamma_c = \{s_1, \ldots, s_t\}$ where each s_i, $1 \leq i \leq t$, is a d-add-remove segment in γ_c. The end-goal of the algorithm is to find an ordered set of vertex labels L such that $\gamma' = label(\gamma, L)$ is nice starting from $S \cap R$. We let γ'_s, $\gamma'_c = \{s'_1, \ldots, s'_t\}$, and γ'_e denote the starting, central, and ending pieces of γ' respectively. Since

Algorithm 3 SOLVE-VCW

Input: A graph G of degree at most d, a vertex cover S of G, and a full unlabeled edit sequence γ (assumed to be d-well-formed).

Output: A full labeled edit sequence $label(\gamma, L)$ which is nice starting from S if one exists and $\varnothing$ otherwise, where L is some ordered set of vertex labels.

1: Let γ_s be the starting piece of γ;
2: Let $\gamma_c = \{s_1, \ldots, s_t\}$ be the central piece of γ;
3: Let γ_e be the ending piece of γ;
4: $S' = \{v \mid v \in S \wedge |N_G(v) \cap \{V(G) \setminus S\}| = 0\}$;
5: Enumerate all independent sets of size $|\gamma_s|$ in S';
6: **for each** independent set
7: Let L be the set of vertex labels (in any order) in the independent set;
8: $\gamma'_s = label(\gamma_s, L)$;
9: **for each** instance $\mathcal{J}_y$, $1 \leq y \leq t$
10: $S'' = V(S, \gamma'_s)$;
11: Create the connected components $\mathcal{CC}_1, \ldots, \mathcal{CC}_y$ (initially empty sets);
12: Generate all mappings between components and d-add-remove segments;
13: **for each** mapping
14: Split γ_c into y pieces $\gamma_1, \ldots, \gamma_y$;
15: $G_A =$ CONST-AUX-GRAPH$(G, S'', \gamma_1, \ldots, \gamma_y)$;
16: Find a multicolored independent set in G_A of size y;
17: **if** (G_A does not contain such a set)
18: **continue**;
19: **else**
20: Let $Q = \{u, \ldots, v\}$ be such a set;
21: **for each** $\lambda \in mix(seq(u), \ldots, seq(v))$
22: Let L be the set of ordered labels in λ;
23: **if** (λ is nice starting from S'' and $\lambda = label(\gamma_c, L)$)
24: $\gamma'_c = label(\gamma_c, L)$;
25: Complete the ending piece γ'_e;
26: $\gamma' = concat(concat(\gamma'_s, \gamma'_c), \gamma'_e)$;
27: **if** (γ' is labeled and nice starting from S)
28: **return** γ';
29: **return** $\varnothing$;

every valid edit sequence can be converted into a nice edit sequence (Lemma 1) and given that the SOLVE-AVCR algorithm tries all possible γ's, we can assume without loss of generality that γ' will be nice. To find γ', we consider each of the pieces of γ separately. Note that if γ' exists, then we know from Definition 3 that the graph induced by $V(s'_i)$, $1 \leq i \leq t$, must be a connected subgraph of $G[R]$. If $V(s'_i)$ and $V(s'_j)$ are separated for all $1 \leq i, j \leq t$ and $i \neq j$, then $G[V(\gamma'_c)]$ is a graph with $t < |\gamma|$ connected components. The number of vertices in and the diameter of each connected component is at most $2d$. However, $V(s'_i)$ and $V(s'_j)$ need not be separated. In fact, it might be the case that $G[V(\gamma'_c)]$ is connected.

The SOLVE-VCW algorithm starts by building the set S', which is equal to the set S' of Algorithm 2 and whose size cannot be "too large" due to line 9 of Algorithm 2. Whenever $|\gamma_s| > 0$, $V(\gamma'_s)$ must be an independent set in S'. Therefore, we enumerate all the possible independent sets of size $|\gamma_s|$ in S'. Each such set will result in a set of $|\gamma_s|$ vertex labels we can assign to the removal markers in γ_s (in any order) to get γ'_s (Algorithm 3, lines 5–8). For each γ'_s, we then attempt to label γ_c. To do so, we generate t instances $\{\mathcal{J}_1, \ldots, \mathcal{J}_t\}$ to cover all possible scenarios in which $G[V(\gamma'_c)]$ has anywhere between 1 and t connected components. Note that we construct the set $S'' = V(S \cap R, \gamma'_s)$ on line 10 of Algorithm 3 so that γ'_c starts from S'' and we can ignore γ'_s for

now. With each instance $\mathcal{J}_y$, $1 \leq y \leq t$, we associate y initially empty sets $\{\mathcal{CC}_1, \ldots, \mathcal{CC}_y\}$. Each set will correspond to a connected component in $G[V(\gamma'_c)]$. Then for each instance $\mathcal{J}_y$, $1 \leq y \leq t$, we enumerate all possible mappings between $\{\mathcal{CC}_1, \ldots, \mathcal{CC}_y\}$ and the segments of γ_c, i.e. every segment of γ'_c will be forced to touch vertices from a single connected subgraph of $G[R]$ which will correspond to some component $\mathcal{CC}_i$ of $G[V(\gamma'_c)]$, $1 \leq i \leq y$. Every connected component $\mathcal{CC}_i$, $1 \leq i \leq y$, will have at most $2dt < |\gamma_c|$ vertices (at most t d-add-remove segments) and therefore its diameter is bounded above by $2dt$. Hence, for every $\mathcal{CC}_i$ there exists a center vertex $v \in V(G[R])$, such that $\mathcal{CC}_i \subseteq B(v, 2dt)$. Lines 13–28 of the Solve-VCW algorithm exploit this fact to construct an auxiliary graph G_A to find a set L of ordered labels such that $\gamma'_c = \mathit{label}(\gamma_c, L)$ is nice starting from S''.

After mapping every segment of γ_c to some component $\mathcal{CC}_i$ in line 12 of Algorithm 3, we split γ_c into y pieces $\gamma_1, \ldots, \gamma_y$, where each piece γ_i is obtained from γ_c by cutting all segments which were not mapped to $\mathcal{CC}_i$, $1 \leq i \leq y$. Each piece corresponds to a d-well-formed full unlabeled edit sequence. Hence, if we call the Const-Aux-Graph algorithm (Algorithm 4) with inputs the graph $G[R]$, the vertex cover S'' (or $V(S \cap R, \gamma'_s)$) of $G[R]$, and the different pieces of γ_c, the resulting auxiliary graph G_A will satisfy Lemma 14. That is, the existence of a multicolored independent set of size y in G_A implies the existence of $L_1, \ldots, L_y$ such that $\gamma'_1 = \mathit{label}(\gamma_1, L_1), \ldots, \gamma'_y = \mathit{label}(\gamma_y, L_y)$ are nice starting from S'', each subgraphs $G[V(\gamma'_1)], \ldots, G[V(\gamma'_y)]$ is connected, and $\gamma'_1, \ldots, \gamma'_y$ are compatible. Therefore, $\gamma'_c \in \mathit{mix}(\gamma'_1, \ldots, \gamma'_y)$ is valid (Definition 2) and nice; as γ'_c consists only of d-add-delete segments all satisfying the connectivity and early removal invariants of Definition 3. On the other hand, if we have a no-instance of the y-MIS problem, then either our mapping from segments to connected components was incorrect or the number of connected components was incorrect. The former case is handled by the fact that we are trying all possible mappings and the latter case is handled by one of the t generated instances.

To complete the ending piece of γ', we simply iterate over $\mathit{concat}(\gamma'_s, \gamma'_c)$ and add an addition marker for every vertex touched an odd number of times in $\mathit{concat}(\gamma'_s, \gamma'_c)$; as γ' must be a walk from $S \cap R$ in $R_{\mathrm{VC}}(G[R], 0, k)$ back to $S \cap R$. If the concatenation of all three sequence γ'_s, γ'_c, and γ'_e produces a nice labeled edit sequence starting from $S \cap R$ then we return it. Otherwise, Algorithm 3 proceeds to the next iteration. If all t instances are no-instances, then we know that the original VCW is a no-instance. □

Lemma 14. *Given a graph G of degree at most d, a vertex cover S of G, and t d-well-formed full unlabeled edit sequences $\gamma_1, \ldots, \gamma_t$ each of size at most ℓ, the* Const-Aux-Graph *algorithm (Algorithm 4) generates a graph G_A with at most t color classes in* **FPT** *time. Each vertex $v \in V(G_A)$ has at most $(d+1)\ell 2^\ell \ell! d^{\ell^2(\ell+1)}$ neighbors in every color class other than $c(v)$. Every color class of G_A has at most $|V(G)| 2^{\ell+1} (d^{\ell+1})^{2\ell+2}$ vertices. Moreover, the existence of a multicolored independent set of size t in G_A implies the existence of $L_1, \ldots, L_t$ such that $\gamma'_1 = \mathit{label}(\gamma_1, L_1), \ldots, \gamma'_t = \mathit{label}(\gamma_t, L_t)$ are nice starting from S, each subgraph $G[V(\gamma'_1)], \ldots, G[V(\gamma'_t)]$ is connected, and $\gamma'_1, \ldots, \gamma'_t$ are compatible.*

Proof. We construct a graph G_A such that a multicolored independent set of size t in G_A implies the existence of $\gamma'_1, \ldots, \gamma'_t$ where:

(1) $\gamma'_1, \ldots, \gamma'_t$ are labeled and full,
(2) $G[V(\gamma'_1)], \ldots, G[V(\gamma'_t)]$ are all connected,
(3) $V(\gamma'_1), \ldots, V(\gamma'_t)$ are all separated, and
(4) $\gamma'_1, \ldots, \gamma'_t$ are all nice starting from S.

Algorithm 4 CONST-AUX-GRAPH

Input: A graph G of degree at most d, a vertex cover S of G, and t full unlabeled edit sequences $\gamma_1, \ldots, \gamma_t$ (assumed to be d-well-formed).
Output: A graph G_A.

```
1: G_A = (V, E);
2: V(G_A) = E(G_A) = ∅;
3: for each γ_i, 1 ≤ i ≤ t
4:    for each v ∈ V(G)
5:       Enumerate every full labeled edit sequence λ of size |γ_i| that only touches vertices in B(v, |γ_i|);
6:       for each λ
7:          Let L be the set of ordered labels in λ;
8:          if (λ is nice starting from S and λ = label(γ_i, L) and G[V(λ)] is connected)
9:             V(G_A) = V(G_A) ∪ {v};
10:            c(v) = c_i (v is assigned color c_i)
11:            Store λ in seq(v);
12: for each (u, v) ∈ V(G_A) such that c(u) ≠ c(v)
13:    if (V(seq(u)) and V(seq(v)) are not separated)
14:       E(G_A) = E(G_A) ∪ {{u, v}} (no double edges);
15: return G_A;
```

When conditions (1), (3), and (4) are satisfied, we know from Observation 5 that $\gamma'_1, \ldots, \gamma'_t$ will be compatible. Conditions (2) and (3) are enforced by our mapping from components to segments in the SOLVE-VCW algorithm (Algorithm 3, lines 11–14) and line 8 of Algorithm 4. If $\gamma'_1, \ldots, \gamma'_t$ satisfying conditions (1) to (4) exist, we know that there must exist some set of distinct vertices $\{u_1, \ldots, u_t\} \in V(G)$ such that $V(\gamma'_1) \subseteq B(u_1, |\gamma'_1|), \ldots, V(\gamma'_t) \subseteq B(u_t, |\gamma'_t|)$.

G_A will contain a vertex assigned color c_i, $1 \leq i \leq t$, for each $C \subseteq V(G)$ satisfying the following properties:

(1) $G[C]$ is connected,
(1) $|C| \in \{|\gamma_1|, \ldots, |\gamma_t|\}$, and
(2) $label(\gamma_i, L)$ is nice starting from S, for L the set of vertex labels in C.

For every γ_i, $1 \leq i \leq t$, we attempt to assign every $v \in V(G)$ the role of a center and enumerate every possible full labeled edit sequence λ of size $|\gamma_i|$ touching vertices in $B(v, |\gamma_i|)$ only. For each λ satisfying the conditions on line 8 of Algorithm 4, we add a vertex in $V(G_A)$ and assign it color c_i. Every vertex in $V(G_A)$ will be associated with a corresponding λ, i.e. for every vertex $u \in V(G_A)$ we let $seq(u) = \lambda$ (Algorithm 4, lines 6–11). After repeating the same process for every γ_i, $V(G_A)$ will have at most t color classes. The edges of G_A are added as follows: For every two vertices $u, v \in V(G_A)$ such that $c(u) \neq c(v)$, there is an edge $\{u, v\}$ if $V(seq(u))$ and $V(seq(v))$ are not separated (Algorithm 4, lines 12–14).

The running time of the CONST-AUX-GRAPH algorithm follows from the description. For each γ_i, $1 \leq i \leq t$, and each $v \in V(G)$, $|B(v, |\gamma_i|)| \leq d^{|\gamma_i|+1} \leq d^{\ell+1}$ by Observation 12. Hence, there are at most $2^{\ell+1}(d^{\ell+1})^{2\ell+2}$ possible full labeled edit sequences λ of size $|\gamma_i| \leq \ell$ to enumerate (Observation 15). Whenever λ is a nice edit sequence starting from S and $G[V(\lambda)]$ is connected, a corresponding vertex u assigned color c_i is added to $V(G_A)$ and $seq(u)$ is set to λ (Algorithm 4, lines 3–11). Hence, every color class of G_A will have at most $|V(G)|2^{\ell+1}(d^{\ell+1})^{2\ell+2}$ vertices.

Note that some color classes in G_A could be empty. To prove the bound on the degree of each vertex, we assume that at least two color classes are non-empty. By Observation 13, each vertex $v \in V(G)$ can appear in at most $d^{\ell^2(\ell+1)}$ vertex-differing connected subgraphs of G of size at

most ℓ. For each of those vertex-differing subgraphs of G, there are at most $\ell!$ ways we can order their corresponding vertex labels. Therefore, there are at most $2^\ell \ell! d^{\ell^2(\ell+1)}$ distinct full labeled edit sequences λ such that λ touches vertex v, $|\lambda| \leq \ell$, and $G[V(\lambda)]$ is connected.

We let $V_i(G_A)$ and $V_j(G_A)$ be any two non-empty color classes in G_A, where $1 \leq i, j \leq t$. We let u be a vertex in $V_i(G_A)$ and we count the number of possible neighbors of u in $V_j(G_A)$. For any vertex $w \in V_j(G_A)$, an edge $\{u, w\}$ in $E(G_A)$ implies that $\{V(seq(u)) \cup N_G(V(seq(u)))\} \cap \{V(seq(w)) \cup N_G(V(seq(w)))\} \neq \emptyset$, as otherwise $V(seq(u))$ and $V(seq(w))$ are separated. Moreover, we know that $G[V(seq(u)) \cup N_G(V(seq(u)))]$ is a connected subgraph of G of size at most $(d+1)|V(seq(u))| \leq (d+1)\ell$. Hence, there are at most $(d+1)\ell 2^\ell \ell! d^{\ell^2(\ell+1)}$ full labeled edit sequences λ such that λ touches some vertex in $V(seq(u)) \cup N_G(V(seq(u)))$, $|\lambda| \leq \ell$, and $G[V(\lambda)]$ is connected. Therefore, there can be at most $(d+1)\ell 2^\ell \ell! d^{\ell^2(\ell+1)}$ vertices $w \in V_j(G_A)$ such that $\{V(seq(u)) \cup N_G(V(seq(u)))\} \cap \{V(seq(w)) \cup N_G(V(seq(w)))\} \neq \emptyset$, as needed. □

Lemma 15. *Every instance of the* Vertex Cover Walk *problem generated in the* Solve-AVCR *algorithm (Algorithm 2) can be solved in* **FPT** *time on graphs of degree at most d.*

Proof. We know from Algorithm 2 (line 8) that the size of S' (Algorithm 3, line 10) is less than $\mathcal{R}(d+2, \ell)$. Hence enumerating all independent sets of fixed size ℓ (or less) in S' can be accomplished in $\mathcal{O}^*(\binom{\mathcal{R}(d+2,\ell)}{\ell})$ time. For each instance $\mathcal{J}_y$, $1 \leq y \leq t \leq \ell$, there are at most y^t possible mappings. So in the worst case, we have $\sum_{y=1}^{t} y^t < t^{t+1} < \ell^{\ell+1}$ mappings to consider. Constructing G_A can be done in **FPT** time (Lemma 14) and the vertices of G_A will satisfy the properties needed to apply Lemma 10 and solve the y-MIS problem in **FPT** time. Every remaining operation can be accomplished in time polynomial in n and ℓ. Therefore, every instance of the VCW problem can be solved in **FPT** time. □

We can finally provide a proof for Theorem 6:

Proof (Theorem 6). Combining Lemmas 12 and 15, we know that there exists an algorithm which solves the VCR problem in **FPT** time on graphs of degree at most d, as needed. □

6 Discussion

To the best of our knowledge, our results constitute the first in-depth study of the VCR problem parameterized by the length of a reconfiguration sequence. We showed that even though the Vertex Cover problem is solvable in polynomial time on bipartite graphs, VCR remains **W**[**1**]-hard. On the tractable side, we showed that VCR is solvable in time polynomial in n for trees as well as cactus graphs and is fixed-parameter tractable for graphs of bounded degree.

Several questions remain unanswered. For instance, it would be interesting to determine the complexity of the VCR problem restricted to 3-regular graphs. Moreover, we believe that our characterization of instances solvable in time polynomial in n can be extended to a wider class of graphs. For example, it is interesting to ask for what graphs is the VCR problem solvable in time polynomial in n if we allow 3-bounded prefixes as opposed to 2-bounded prefixes for cactus graphs and 1-bounded prefixes for trees. Moreover, our polynomial-time algorithms imply a bound on the diameter of the reconfiguration graph which is linear in n. In particular, we showed that for both trees and cactus graphs the diameter of the reconfiguration graph is an additive constant away from the size of the largest symmetric difference of any two vertex covers of the input graph. We feel that this relationship between tractability and symmetric difference deserves a more careful study.

Finally, we believe that the techniques used in both our hardness proofs and positive results can be extended to cover a host of graph deletion problems defined in terms of hereditary graph properties [21]. It also remains to be seen whether our **FPT** result can be extended to a larger class of sparse graphs similar to the work of Fellows et al. on local search [8].

References

1. Abu-Khzam, F. N., Collins, R. L., Fellows, M. R., Langston, M. A., Suters, W. H., and Symons, C. T. Kernelization algorithms for the vertex cover problem: Theory and experiments. In *Proc. of the Sixth Workshop on Algorithm Engineering and Experiments* (2004), pp. 62–69.
2. Bonsma, P. The complexity of rerouting shortest paths. In *Proc. of Mathematical Foundations of Computer Science* (2012), pp. 222–233.
3. Cereceda, L., van den Heuvel, J., and Johnson, M. Connectedness of the graph of vertex-colourings. *Discrete Mathematics 308*, 56 (2008), 913–919.
4. Chor, B., Fellows, M. R., and Juedes, D. Linear kernels in linear time, or how to save k colors in $O(n^2)$ steps. In *Proc. of the 30th International Conference on Graph-Theoretic Concepts in Computer Science* (2004), pp. 257–269.
5. Diestel, R. *Graph Theory*. Springer-Verlag, Electronic Edition, 2005.
6. Downey, R. G., and Fellows, M. R. *Parameterized complexity*. Springer-Verlag, New York, 1997.
7. Fellows, M. R., Hermelin, D., Rosamond, F., and Vialette, S. On the parameterized complexity of multiple-interval graph problems. *Theor. Comput. Sci. 410*, 1 (2009), 53–61.
8. Fellows, M. R., Rosamond, F. A., Fomin, F. V., Lokshtanov, D., Saurabh, S., and Villanger, Y. Local search: is brute-force avoidable? In *Proc. of the 21st International Joint Conference on Artifical Intelligence* (2009), pp. 486–491.
9. Flum, J., and Grohe, M. *Parameterized complexity theory*. Springer-Verlag, Berlin, 2006.
10. Fricke, G., Hedetniemi, S. M., Hedetniemi, S. T., and Hutson, K. R. γ-Graphs of Graphs. *Discussiones Mathematicae Graph Theory 31*, 3 (2011), 517–531.
11. Garey, M. R., Johnson, D. S., and Stockmeyer, L. J. Some simplified NP-complete graph problems. *Theor. Comput. Sci. 1*, 3 (1976), 237–267.
12. Gopalan, P., Kolaitis, P. G., Maneva, E. N., and Papadimitriou, C. H. The connectivity of boolean satisfiability: computational and structural dichotomies. *SIAM J. Comput. 38*, 6 (2009), 2330–2355.
13. Hall, P. On representatives of subsets. In *Classic Papers in Combinatorics*. Birkhuser Boston, 1987, pp. 58–62.
14. Harary, F. *Graph theory*. Addison-Wesley, 1991.
15. Ito, T., Demaine, E. D., Harvey, N. J. A., Papadimitriou, C. H., Sideri, M., Uehara, R., and Uno, Y. On the complexity of reconfiguration problems. *Theor. Comput. Sci. 412*, 12-14 (2011), 1054–1065.
16. Ito, T., Kamiński, M., and Demaine, E. D. Reconfiguration of list edge-colorings in a graph. *Discrete Applied Mathematics 160*, 15 (2012), 2199–2207.
17. Ito, T., Kawamura, K., Ono, H., and Zhou, X. Reconfiguration of list L(2, 1)-labelings in a graph. In *Proc. of the 23rd Annual International Symposium on Algorithms and Computation* (2012), pp. 34–43.
18. Kamiński, M., Medvedev, P., and Milanič, M. Shortest paths between shortest paths. *Theor. Comput. Sci. 412*, 39 (2011), 5205–5210.
19. Kamiński, M., Medvedev, P., and Milanič, M. Complexity of independent set reconfigurability problems. *Theor. Comput. Sci. 439* (June 2012), 9–15.
20. Khot, S., and Raman, V. Parameterized complexity of finding subgraphs with hereditary properties. *Theor. Comput. Sci. 289*, 2 (2002), 997–1008.
21. Mouawad, A. E., Nishimura, N., Raman, V., Simjour, N., and Suzuki, A. On the parametereized complexity of reconfiguration problems. In *Proc. of the 8th International Symposium on Parameterized and Exact Computation* (2013).
22. Niedermeier, R. *Invitation to fixed-parameter algorithms*. Oxford University Press, 2006.